
tree-gen Documentation

QuTech, TU Delft

Sep 22, 2023

1	Overview	3
1.1	Base classes and primitives	3
1.2	Defining node types	5
1.3	Directives	6
1.4	Generated APIs	7
2	Directory tree example	13
2.1	Python usage	22
2.2	Complete file listings	28
3	Interpreter example	41
3.1	External nodes & multiple trees	41
3.2	Integration with Flex/Bison (or other parsers)	41
3.3	Visitor pattern	45
3.4	Complete file listings	46

tree-gen is a program that generates C++ source code for tree structures based on a simple and concise input language. The API documentation is [generated by Doxygen](#).

tree-gen is a code generator that outputs all the repetitive stuff needed for defining a proper tree structure in C++.

Trees are described by *.tree files. Such a file consists of a number of directives at the top and one or more tree node descriptions. Each tree node gets its own C++ class, with a bunch of methods defined on it for traversal, the visitor pattern, safe typecasting, cloning, and so on. It would be extremely tedious and error prone to define all this stuff manually, hence the generator.

It is important to realize and perhaps not immediately obvious that the recursive structure in a *.tree file represents C++ inheritance rather than the actual tree structure. You may for instance encounter something like this:

```
expression {
  addition {
    ...
  }
  subtraction {
    ...
  }
}
```

That doesn't mean an expression consists of an addition and a subtraction, but rather that an addition is a type of expression, and subtraction is another.

1.1 Base classes and primitives

tree-gen trees consist of three kinds of objects: nodes, edges, and primitives.

The nodes are classes generated by tree-gen based on some base class. This is typically tree::base::Base. The class inheritance tree of the node types is used for “union” types, i.e. to allow some node in the tree to be one of a number of types. For example, some node may want to have an expression child node, which in practice would be for instance an addition node or a subtraction node. At any level of this class hierarchy, nodes can be given attributes to store data. These attributes are either edges or primitives.

Edges define the relation between some parent node and some number of child nodes. Six edge classes are defined, resulting in different relations between the nodes.

- `tree::base::Maybe<T>`: the parent will own zero or one child node of (super)type T.
- `tree::base::One<T>`: the parent will logically own exactly one child node of (super)type T.
- `tree::base::Any<T>`: the parent will own zero or more child nodes of (super)type T.
- `tree::base::Many<T>`: the parent will logically own one or more child node of (super)type T.
- `tree::base::OptLink<T>`: the parent links to zero or one child node of (super)type T elsewhere in the tree.
- `tree::base::Link<T>`: the parent logically links to exactly one child node of (super)type T elsewhere in the tree.

The “logically” term used above alludes to the fact that One/Many/Link don’t necessarily need to have at least one child at all times. They may for instance be empty during construction. But at least one child node is required for the tree to be considered to be “complete”. Furthermore, completeness with respect to some root node requires that any Link or non-null OptLink refers to a node that is actually reachable from that root node.

Maybe and One are based on `std::shared_ptr`, Any and Many are based on `std::vector<One>`, and OptLink and Link are based on `std::weak_ptr`, but these types are abstracted away from the user almost entirely. Notably, all exposed dereference operations are null- and range-checked, throwing exceptions if there’s a problem, rather than causing a segmentation fault down the line.

The namespace for the Base and edge classes can be set using the `tree_namespace` directive in case you want to override them or add functionality. Note that you can use “using One = tree::base::One;” etc. if you only want to override part of the classes. You can also opt to use the `tree-all.[ch]pp.inc` files, which define the same structures that exist in the `tree` namespace, but allows you to override some of the standard library types they are based on, for example to use a different base exception type. If the `tree_namespace` directive is not specified, the namespace that the node classes are generated into will be used by default.

Finally, primitives are used to store the actual data, forming the leaves of the tree. They can generally be any externally-defined C++ type, although a few template functions need to be defined and specialized for the types you want to use. These are:

```
template <typename T>
T initialize() { ... }

template <typename T>
void serialize(const T &obj, tree::cbor::MapWriter &map) { ... }

template <typename T>
T deserialize(const tree::cbor::MapReader &map) { ... }
```

The `initialize` function is used to construct a primitive of type T in such a way that it has a defined value. This is in contrast to the behavior of C’s own primitive types (`int`, `bool`, etc.), which are specified to be undefined until assigned. This is avoided using this `initialize` function to prevent confusion caused by undefined behavior. Typically, `initialize()` can be implemented using just `return T{};` for the general case to defer to the constructor for complex types, and then be specialized for the primitive types that you want to use.

The `serialize` and `deserialize` functions are optional. When given, logic is generated to serialize and deserialize entire trees to and from a CBOR representation.

In addition to the type-safe primitives declared in the tree description file, nodes can also be annotated using arbitrary types. Specifically, every node can be annotated with zero or one annotation of every C++ type. So, any struct or class you make can be attached to any node in the tree, regardless of the tree specification. This is useful, for instance, to carry metadata, such as line number or debugging information. When `tree-gen` is used to generate some user-facing

tree structure in a library, users of the library can also use this to add their own data to a tree as they operate on it, to prevent them from having to convert to their own data structure first.

Note that the above requires RTTI to be enabled and supported by the compiler, but there's really no good reason to disable that nowadays.

1.2 Defining node types

Nodes have the following form in the tree file:

```
# [documentation for node class]
[snake_case_node_name] {

    # [documentation for primitive]
    [snake_case_member_name]: [C++ namespace path];

    # [documentation for edge/child node]
    [snake_case_member_name]: [Maybe|One|Any|Many|OptLink|Link]<[snake_case_node_
↪name]>;

    # [documentation for external edge/child node]
    [snake_case_member_name]: external [Maybe|One|Any|Many|OptLink|Link]<[C++_
↪namespace path]>;

    [zero or more specializations of this node; recursive structure]

    [optional reorder directive]
}
```

Nodes can have zero or more primitives, zero or more edges/child nodes, zero or more *external* edges/child nodes, and zero or more specializations.

The difference between normal edges and external edges is that the type for a normal edge must resolve to another node in the same tree file, while an external edge may use any C++ type name for the node class. That allows you to refer to nodes defined in other tree files as well. Note however that the generator expects external nodes to actually be node classes, such that it can properly generate the debug dump indentation logic, among other things.

Note that while node type names are specified in `snake_case` in the tree file, the name will be converted to `TitleCase` for the generated class names to follow decent C++ naming conventions. This is just because it's easier to convert snake case to title case than vice versa, and the generator uses both forms (title case for the class names, snake case for methods and members).

Note also the comments explicitly added to the example. Comments using a `#` are interpreted as docstrings by the generator; they are copied into the source code using javadoc-style comment blocks. This also means you can't use `#`-"comments" everywhere, since they're actually a grammatical construct. Therefore, tree-gen also supports `//` comments. These are completely ignored by the parser.

The order in which nodes are defined doesn't matter. This allows you to make recursive tree structures, without having to worry about forward declarations. The generator handles those. The order in which the attributes of a node are defined does matter, though, as it's used for the order in which the values can be passed to the node constructor; they are ordered subclass to superclass, top to bottom.

The order of the primitives and child nodes (a.k.a. fields) is used in a few places, primarily the constructor and for debug dump output generation. The order is as specified within each class hierarchy level, but the fields of the more-specialized classes come before the fields of their ancestors.

When additions are made to a tree, it may be required to maintain compatibility with older tree versions by maintaining

field order. In order to support this, the above field order can be overridden using a `reorder` directive. This directive has the following syntax:

```
# [optional documentation (ignored)]
reorder(a, b, c);
```

where `a`, `b`, `c` is the desired order of the fields. Any fields not explicitly specified in the list will automatically appear at the end, using the default order.

1.3 Directives

The following directives exist. They should be placed at the top of the tree file, before the first node. Their order doesn't matter, unless otherwise specified. Note that most directives are required.

- `source`: used to specify documentation for the generated source file; any docstring above the directive is copied into the file as file-level doxygen documentation.
- `header`: used to specify documentation for the generated header file; any docstring above the directive is copied into the file as file-level doxygen documentation.
- `python`: used to specify documentation for the generated Python file (if enabled via the command line); any docstring above the directive is copied into the Python file's docstring.
- `tree_namespace` `<namespace::path>`: the namespace that the Base and edge classes live in. You should set this to `tree::base` unless you intend to override or use different base classes. This is unused in Python.
- `initialize_function` `<namespace::path::initialize>`: the name (including namespace path leading up to it) of the `T()` function used for getting the default value of any of the used primitive classes. Unused in Python; here, the default constructor is used regardless of type (as it always exists and is meaningful in Python).
- `serdes_functions` `<namespace::path::serialize>` `<namespace::path::deserialize>`: optionally, the names (including namespace paths leading up to them) of the functions used to respectively serialize and deserialize primitive classes. If not specified, serialization functionality is disabled in C++. Python always has serdes support, but support is augmented when a serdes function pair is specified. The Python generator interprets the namespace path as a module path. Refer to the serialization section for more info.
- `location` `<namespace_path::SourceLocation>`: optionally, the name of the source location annotation class (including namespace path leading up to it). The debug dump will look for an annotation of this type on each node when doing a debug dump, and if it exists, uses it to add source information to the node, by streaming out a # followed by the stream overload for the class. The base class needs to be capable of annotations if you use this. This is not supported in Python.
- `include` `"<path>"`: adds an `#include` statement to the top of the generated C++ header file.
- `src_include` `"<path>"`: like `include`, but adds to the top of the generated C++ source file only.
- `import ... , from ... import ...`: adds Python import statements to the top of the generated module.
- `namespace` `<name>`: used to specify the namespace for the generated tree classes. As in C++, you need multiple of these to specify the full path. The docstring for the *first* annotation of this type that has a docstring in front is used to document the innermost namespace javadoc-style for Doxygen documentation. This is not used in Python.

1.4 Generated APIs

The following methods are generated for each node class:

- a constructor with all members of the node in its signature as optional values, defaulting to the values returned by the initializer function.
- `bool is_complete() const`: returns whether the node its called on and any subtree rooted in that node is fully defined (no empty One, Many or Link, all array entries in Any/Many nonempty, and all Links and defined OptLinks reachable from the root node).
- `%NodeType type() const`: returns the type of this node, using the also-generated `%NodeType` enumeration.
- `One<Node> copy() const`: returns a shallow copy of this node.
- `One<Node> clone() const`: returns a deep copy of this node.
- A value-based equality operator. Note that this ignores equality of any annotations.
- `void visit(Visitor &visitor)` (C++ only): implements the visitor pattern using the also-generated abstract Visitor/RecursiveVisitor classes. See below.
- `void dump(std::ostream &out=std::cout, int indent=0)`: does a debug dump of the node to the given stream with the given indentation level.
- `SomeNodeType *as_some_node_type()` (C++ only): does the equivalent of a `dynamic_cast` to the given node type, returning `this` if the type is correct or `nullptr` if not.

Note that the prototypes for the Python equivalents differ as appropriate.

An implicit node class simply named `Node` is always generated, serving as the base class for all other nodes. In C++, it is what derives from the `Base` class defined in the namespace specified using `tree_namespace`. In Python, the `Node` class always derives directly from `object`.

1.4.1 Tree traversal

Tree traversal is accomplished by starting at the root and working your way down. The attributes specified in the tree file appear as public members of the node classes. Maybe, One, OptLink, and Link can be dereferenced simply using the `*` or `->` operators; Any and Many use the `[]` indexation operator. Note that these are all null- and range-checked, unlike the `stdlib` equivalents.

Tree traversal is complicated by the fact that you often don't know exactly what the type of a node is. For example, an expression node may actually be a subtraction or addition. For this, tree-gen supports a few different patterns in C++:

- *Visitor pattern.* You define a class inheriting from the generated `Visitor` or `RecursiveVisitor` classes. These abstract classes provide or require you to implement functions for each node type. The appropriate one then gets called when you pass an instance of your visitor class to the `visit()` method on a node. You must always override `visit_node()`, usually to throw a suitable exception in case an unexpected node is encountered. The default implementation for all the other node types differs between the two classes. For `Visitor`, the default implementation falls back to the next more generically typed function (`visit_subtraction()` falls back to `visit_expression()`, `visit_expression()` falls back to `visit_node()`). For `RecursiveVisitor`, the default implementation for non-leaf node types recursively calls `visit()` for all child nodes, thus recursively traversing the tree. For leaf nodes, both visitor classes have the same behavior.
- *Using the 'as_()' methods.** Given for instance an expression node, you might do

```

if (auto addition = expression.as_addition()) {
    ...
} else if (auto subtraction = expression.as_subtraction()) {
    ...
} else {
    ...
}

```

Because a `nullptr` evaluates to false in C++, the blocks will only be executed if the cast succeeds. Be careful copy-pasting though; you can accidentally use the `addition` variable in the `subtraction` block if you'd want, but it's obviously null in that case. C++ scoping is weird.

- *Using a switch statement.* You might do

```

switch (expression.type()) {
    case NodeType::Addition:
        ...
    case NodeType::Subtraction:
        ...
}

```

This doesn't handle the cast for you, but in cases where you only need to switch based on the type and don't need access to members of the nodes this is more descriptive than the if/else form.

Just choose the method that makes the most sense within context. Python is so much more dynamic in general that you're better off using its duck typing functionality or using `isinstance()` directly, so no special features are generated for this.

Note that tree-gen trees do *not* contain allow traversal back toward the root of a tree. Supporting this would greatly complicate the internals and the user-facing API, because then you wouldn't be able to just move nodes around. It is therefore key that you design your trees and interfaces such that this information is not needed. If this is somehow impossible, you'll have to manage the links back up the tree manually using `(Opt)Link` edges.

1.4.2 Serialization and deserialization

Optionally, logic to serialize and deserialize trees can be generated in addition to the APIs above. To enable this, you have to tell tree-gen how to serialize and deserialize the primitive types that you use in the tree through two templated functions that must be specialized for all primitive types (similar and in addition to the initialization function):

```

template <typename T>
void serialize(const T &obj, tree::cbor::MapWriter &map) { ... }

template <typename T>
T deserialize(const tree::cbor::MapReader &map) { ... }

```

The former must serialize `obj` by calling the various `append_*()` functions on the given `tree::cbor::MapWriter` object. The latter must perform the reverse operation. The namespace(s) and names of these functions must be provided to tree-gen using the optional `serdes_functions` directive. Specifying this directive enables the serialization and deserialization logic.

Once enabled, the entry point for serializing a tree is `tree::base::serialize()` or `tree::base::serialize_file()`, and deserializing is `tree::base::deserialize()` or `tree::base::deserialize_file()`. The internal serialization and deserialization functions in the edge classes and generated node classes are public and could also be used, but these require more boilerplate due to link handling.

If you attempt to use these methods for a tree without serialization/deserialization support enabled, you'll receive a template error. Template errors are notoriously difficult to understand, so you have been warned. Don't use them

unless you enable support through the `serdes_functions` directive. Attempting to serialize a tree that is not well-formed will lead to a `tree::base::NotWellFormed` exception.

Serialization format

The serialization format makes use of the RFC7049 CBOR data representation, serving as a compromise between the readability of JSON and speed/simplicity of the serialization and deserialization logic. Specifically, CBOR can be losslessly converted to and from JSON using third-party tools if need be for debugging (at least for as far as it's used here), but is itself a simple binary format.

Each edge and each primitive receives its own map in the tree. Note that Any and Many are internally represented as a vector of One edges; therefore, Any/Many results in two nested edge objects. The data for nodes and their annotations is stored along with the One/Maybe edges. The keys for the maps are the following:

```
empty Maybe:
  {
    "@T": "?",
    "@t": null
  }

filled Maybe:
  {
    "@T": "?",
    "@i": <sequence number>,
    "@t": "<TitleCase node type name>",
    "<snake_case_attribute_name>": { <attribute data> },
    ...
    "{<annotation type>}": { <annotation data> },
    ...
  }

One:
  {
    "@T": "1",
    "@i": <sequence number>,
    "@t": "<TitleCase node type name>",
    "<snake_case_attribute_name>": { <attribute data> },
    ...
    "{<annotation type>}": { <annotation data> },
    ...
  }

Any:
  {
    "@T": "*",
    "@d": [
      <`One` object for each item>
    ]
  }

Many:
  {
```

(continues on next page)

```

    "@T": "+",
    "@d": [
        <`One` object for each item>
    ]
}

empty OptLink:
{
    "@T": "@",
    "@l": null
}

filled OptLink:
{
    "@T": "@",
    "@l": <sequence number of linked node>
}

Link:
{
    "@T": "$",
    "@l": <sequence number of linked node>
}

```

The @T entry stores which type of edge an object is. This is not strictly needed when deserializing as this information should be known by context; instead, it is used as a validity check. The @i sequence numbers are unique integers for all One/Maybe edges in the tree, which are used to recover the (Opt)Link pointers through their @l key. The @t key is used to recover subtype information, as the tree only knows the abstract supertype of a node in general. These are, however, not necessarily globally unique; they only need to distinguish between subclasses, and therefore just the TitleCase name of the node type without C++ namespace is sufficient.

Keys that start with a { and close with a } are used for annotations. The string enclosed within the {} in the key is used to store the annotation type. The identifier used can either be generated automatically by C++ using typeid(T).name(), or can be specified manually to have more control. Any annotation type that doesn't have a serialization and deserialization function registered for it within tree::annotatable::serdes_registry is silently ignored in either direction.

All remaining keys map to the node attributes using their snake_case name. The corresponding value is (recursively) one of the structures above for edges, or a map containing user-specified key/value pairs for primitives, serialized and deserialized using the functions specified by the serdes_functions directive.

The names of the keys are chosen such that, when ordered by ASCII value, the order is as specified, and such that there can never be name conflicts. The commonly used keys and values are short to minimize serialization and deserialization overhead.

1.4.3 Python support

In addition to C++, tree-gen can also generate pure-Python objects to represent the tree. As in C++, each node type becomes a class, and class hierarchy is used similarly. The edge classes don't exist, however; instead, their functionality is baked into the generated code to prevent an unnecessary user-facing level of indirection during tree traversal that cannot be hidden in Python. Despite this, the generated structures provide a similar level of type safety as the C++ structures do; all writes to fields of node classes are type-checked.

Where applicable, the generated Python file assumes that the Python module tree corresponds with the namespace hierarchy in the C++ world. This applies to the following things:

- primitive type declarations;
- external node types;
- (de)serialization functions.

The initialization function isn't used in Python; rather, the default constructor of the primitive types is used.

The connection between the C++ and Python world is handled through CBOR serialization and deserialization rather than providing Python wrappers around the C++ objects. While this approach isn't as performant as the alternative, it allows a much simpler and more Pythonic interface to be exposed to the user. Serdes is therefore perhaps the most important feature of the Python generator. Note however that it is entirely up to the program or library using tree-gen to provide the requisite interface between the C++ and Python worlds (for example through swig); tree-gen just ensures that the Python and C++ tree implementations always remain in sync.

The serialization and deserialization functions naturally have a different signature in Python than they do in C++. More specifically, the functions must look like this:

```
def serialize(typ, val):
    # typ is either a type (for primitives) or a str (for annotations). In
    # the former case, isinstance(val, typ) may be assumed. Return value must
    # be an object consisting of only the following types:
    # - `dict`s with `str` keys;
    # - `list`s;
    # - `int`s between  $-2^{63}$  and  $2^{63}-1$  inclusive;
    # - `float`s;
    # - `True`, `False`, or `None`;
    # - `str`s;
    # - `bytes` objects.
    # The toplevel object must be a dict.
    pass # TODO

def deserialize(typ, val):
    # typ is either a type (for primitives) or a str (for annotations). val
    # is a structure representing the CBOR serialization using the same
    # Python primitives that can be returned by serialize(). The return value
    # must be an instance of the specified `typ`e for primitives, but can be
    # anything for annotations; for unknown annotation types, just returning
    # val as-is is the recommended fallback behavior.
    pass # TODO
```

Primitive serialization and deserialization can also be done using class methods:

```
class SomePrimitive:
    def serialize_cbor(self):
        # Called instead of serialize(SomePrimitive, self) if it exists.
        pass # TODO

    @staticmethod
    def deserialize_cbor(cbor):
        # Called instead of deserialize(SomePrimitive, cbor) if it exists.
        pass # TODO
```

If no `serdes_function` directive is specified, serdes logic is still generated on the Python end; in this case, only the class method variant is used, and annotations must be directly serializable to CBOR for them to be serialized (they will be silently ignored if they're not).

Directory tree example

This example illustrates the tree system with a Windows-like directory tree structure. The `System` root node consists of one or more drives, each of which has a drive letter and a root directory with named files and subdirectories. To make things a little more interesting, a symlink-like “mount” is added as a file type, which links to another directory.

Using this tree, this example should teach you the basics of using tree-gen trees. The tutorial runs you through the C++ code of `main.cpp` and Python code of `main.py`, but be sure to check out `directory.tree`, `primitives.hpp`, `primitives.hpp`, and (if you want to reproduce it in your own project) `CMakeLists.txt` copied to the bottom of this page as well. You will also find the complete `main.cpp` and `main.py` there.

Let’s start by making the root node using `tree::base::make()`. This works somewhat like `std::make_shared()`, but instead of returning a `shared_ptr` it returns a `One` edge. This might come off as a bit odd, considering trees in graph theory start with a node rather than an edge. This is just a choice, though; a side effect of how the internal “`shared_ptr`”s work and how `Link`/`OptLink` edges work (if you’d store the tree as the root structure directly, the root node wouldn’t always be stored in the same place on the heap, breaking link nodes).

```
auto system = tree::base::make<directory::System>();
```

At all times, you can use the `dump()` method on a node to see what it looks like for debugging purposes. It looks like this:

```
system->dump();
```

↓

```
System(  
  drives: !MISSING  
)
```

While this system node exists as a tree in memory and tree-gen seems happy, our system tree is at this time not “well-formed”. A tree node (or an edge to one) is only considered well-formed if all of the following are true:

- All `One` edges in the tree connect to exactly one node.
- All `Many` edges in the tree connect to at least one node.
- All `Link` and non-empty `OptLink` nodes refer to a node reachable from the root node.

- Each node in the tree is only used by a non-link node once.

Currently, the second requirement is not met.

```
ASSERT(!system.is_well_formed());
```

You can get slightly more information using `check_well_formed()`; instead of returning a boolean, it will throw a `NotWellFormed` exception with an error message.

```
ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
```

↓

```
tree::base::NotWellFormed exception message: 'Many' edge of type N9directory5DriveE_
↳is empty
```

Note that the name of the type is “mangled”. The exact output will vary from compiler to compiler, or even from project to project. But hopefully it’ll be readable enough to make sense of in general.

To fix that, let’s add a default drive node to the tree. This should get drive letter A, because `primitives::initialize()` is specialized to return that for “Letter”s (see `primitives.hpp`).

```
system->drives.add(tree::base::make<directory::Drive>());
```

`Drive` has a `One` edge that is not empty, though, so the tree still isn’t well-formed. Let’s add one of those as well.

```
system->drives[0]->root_dir = tree::base::make<directory::Directory>();
```

Now we have a well-formed tree. Let’s have a look:

```
system.check_well_formed();
system->dump();
```

↓

```
System(
  drives: [
    Drive(
      letter: A
      root_dir: <
        Directory(
          entries: []
          name:
        )
      >
    )
  ]
)
```

We can just change the drive letter by assignment, as you would expect.

```
system->drives[0]->letter = 'C';
```

Before we add files and directories, let’s make a shorthand variable for the root directory. Because `root_dir` is an edge to another node rather than the node itself, and thus acts like a pointer or reference to it, we can just copy it into a variable and modify the variable to update the tree. Note that you’d normally just write “auto” for the type for brevity; the type is shown here to make explicit what it turns into.

```
directory::One<directory::Directory> root = system->drives[0]->root_dir;
```

Let's make a "Program Files" subdirectory and add it to the root.

```
auto programs = tree::base::make<directory::Directory>(
    tree::base::Any<directory::Entry>{},
    "Program Files");
root->entries.add(programs);
system.check_well_formed();
```

That's quite verbose. But in most cases it can be written way shorter. Here's the same with the less versatile but also less verbose `emplace()` method (which avoids the `tree::make()` call, but doesn't allow an insertion index to be specified) and with a "using namespace" for the tree. `emplace()` can also be chained, allowing multiple files and directories to be added at once in this case.

```
{
    using namespace directory;

    root->entries.emplace<Directory>(Any<Entry>{}, "Windows")
        .emplace<Directory>(Any<Entry>{}, "Users")
        .emplace<File>("lots of hibernation data", "hiberfil.sys")
        .emplace<File>("lots of page file data", "pagefile.sys")
        .emplace<File>("lots of swap data", "swapfile.sys");
}
system.check_well_formed();
```

In order to look for a file in a directory, you'll have to make your own function to iterate over them. After all, tree-gen doesn't know that the name field is a key; it has no concept of a key-value store. This is simple enough to make, but to prevent this example from getting out of hand we'll just use indices for now.

Let's try to read the "lots of data" string from `pagefile.sys`.

```
ASSERT(root->entries[4]->name == "pagefile.sys");
// ASSERT(root->entries[4]->contents == "lots of page file data");
// '-> No member named 'contents' in 'directory::Entry'
```

Hm, that didn't work so well, because C++ doesn't know that entry 4 happens to be a file rather than a directory or a mount. We have to tell it to cast to a file first (which throws a `std::bad_cast` if it's not actually a file). The easiest way to do that is like this:

```
ASSERT(root->entries[4]->as_file()->contents == "lots of page file data");
```

No verbose casts required; tree-gen will make member functions for all the possible subtypes.

While it's possible to put the same node in a tree twice (without using a link), this is not allowed. This isn't checked until a well-formedness check is performed, however (and in fact can't be without having access to the root node).

```
root->entries.add(root->entries[0]);
ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
```

↓

```
tree::base::NotWellFormed exception message: Duplicate node of type_
↳N9directory5EntryEat address 0x5588663ab440 found in tree
```

Note that we can index nodes Python-style with negative integers for `add()` and `remove()`, so `remove(-1)` will remove the broken node we just added. Note that the `-1` is not actually necessary, though, as it is the default.

```
root->entries.remove(-1);
system.check_well_formed();
```

We *can*, of course, add copies of nodes. That's what copy (shallow) and clone (deep) are for. Usually you'll want a deep copy, but in this case shallow is fine, because a File has no child nodes. Note that, even for a deep copy, links are not modified; it is intended to copy a subtree to another part of the same tree. To make a complete copy of a tree that maintains link semantics (i.e. does not link back to the original tree) the best way would be to go through serialize/deserialize.

```
root->entries.add(root->entries[0]->copy());
system.check_well_formed();
```

Note that the generated classes don't care that there are now two directories named Program Files in the root. As far as they're concerned, they're two different directories with the same name. Let's remove it again, though.

```
root->entries.remove();
```

Something we haven't looked at yet are links. Links are edges in the tree that, well, turn it into something that isn't strictly a tree anymore. While One/Maybe/Any/Many require that nodes are unique, Link/OptLink require that they are *not* unique, and are present elsewhere in the tree. Let's make a new directory, and mount it in the Users directory.

```
auto user_dir = tree::base::make<directory::Directory>(
    tree::base::Any<directory::Entry>{},
    "");
root->entries.emplace<directory::Mount>(user_dir, "SomeUser");
```

Note that user_dir is not yet part of the tree. emplace works simply because it doesn't check whether the directory is in the tree yet. But the tree is no longer well-formed now.

```
ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
```

↓

```
tree::base::NotWellFormed exception message: Link to node of type_
↳N9directory9DirectoryE at address 0x5588663abcc0 not found in tree
```

To make it work again, we can add it as a root directory to a second drive.

```
system->drives.emplace<directory::Drive>('D', user_dir);
system.check_well_formed();
```

A good way to confuse a filesystem is to make loops. tree-gen is okay with this, though.

```
system->drives[1]->root_dir->entries.emplace<directory::Mount>(root, "evil link to C:
↳");
system.check_well_formed();
```

The only place where it matters is in the dump function, which only goes one level deep. After that, it'll just print an ellipsis.

```
system->dump();
```

↓

```
System(
  drives: [
```

(continues on next page)

(continued from previous page)

```
Drive(
  letter: C
  root_dir: <
    Directory(
      entries: [
        Directory(
          entries: []
          name: Program Files
        )
        Directory(
          entries: []
          name: Windows
        )
        Directory(
          entries: []
          name: Users
        )
        File(
          contents: lots of hibernation data
          name: hiberfil.sys
        )
        File(
          contents: lots of page file data
          name: pagefile.sys
        )
        File(
          contents: lots of swap data
          name: swapfile.sys
        )
        Mount(
          target --> <
            Directory(
              entries: [
                Mount(
                  target --> <
                    ...
                  >
                  name: evil link to C:
                )
              ]
              name:
            )
          >
          name: SomeUser
        )
      ]
      name:
    )
  >
)
Drive(
  letter: D
  root_dir: <
    Directory(
      entries: [
        Mount(
          target --> <
```

(continues on next page)

```
Directory(
  entries: [
    Directory(
      entries: []
      name: Program Files
    )
    Directory(
      entries: []
      name: Windows
    )
    Directory(
      entries: []
      name: Users
    )
    File(
      contents: lots of hibernation data
      name: hiberfil.sys
    )
    File(
      contents: lots of page file data
      name: pagefile.sys
    )
    File(
      contents: lots of swap data
      name: swapfile.sys
    )
    Mount(
      target --> <
        ...
      >
      name: SomeUser
    )
  ]
  name:
)
>
name: evil link to C:
)
]
name:
)
>
)
]
)
```

Now that we have a nice tree, let's try the serialization and deserialization functionality. Serializing is easy:

```
std::string cbor = tree::base::serialize(system);
```

Let's write it to a file; we'll load this in Python later.

```
{
  std::ofstream cbor_output;
  cbor_output.open("tree.cbor", std::ios::out | std::ios::trunc | std::ios::binary);
  cbor_output << cbor;
}
```

Let's also have a look at a hexdump of that.

```
int count = 0;
for (char c : cbor) {
    if (count == 16) {
        std::printf("\n");
        count = 0;
    } else if (count > 0 && count % 4 == 0) {
        std::printf(" ");
    }
    std::printf("%02X ", (uint8_t)c);
    count++;
}
std::printf("\n");
```

↓

```
BF 62 40 54 61 3F 62 40 69 00 62 40 74 66 53 79
73 74 65 6D 66 64 72 69 76 65 73 BF 62 40 54 61
2B 62 40 64 9F BF 62 40 54 61 31 62 40 69 01 62
40 74 65 44 72 69 76 65 66 6C 65 74 74 65 72 BF
63 76 61 6C 18 43 FF 68 72 6F 6F 74 5F 64 69 72
BF 62 40 54 61 31 62 40 69 02 62 40 74 69 44 69
72 65 63 74 6F 72 79 67 65 6E 74 72 69 65 73 BF
62 40 54 61 2A 62 40 64 9F BF 62 40 54 61 31 62
40 69 03 62 40 74 69 44 69 72 65 63 74 6F 72 79
67 65 6E 74 72 69 65 73 BF 62 40 54 61 2A 62 40
64 9F FF FF 64 6E 61 6D 65 BF 63 76 61 6C 6D 50
72 6F 67 72 61 6D 20 46 69 6C 65 73 FF FF BF 62
40 54 61 31 62 40 69 04 62 40 74 69 44 69 72 65
63 74 6F 72 79 67 65 6E 74 72 69 65 73 BF 62 40
54 61 2A 62 40 64 9F FF FF 64 6E 61 6D 6D 65 BF 63
76 61 6C 67 57 69 6E 64 6F 77 73 FF FF BF 62 40
54 61 31 62 40 69 05 62 40 74 69 44 69 72 65 63
74 6F 72 79 67 65 6E 74 72 69 65 73 BF 62 40 54
61 2A 62 40 64 9F FF FF 64 6E 61 6D 65 BF 63 76
61 6C 65 55 73 65 72 73 FF FF BF 62 40 54 61 31
62 40 69 06 62 40 74 64 46 69 6C 65 68 63 6F 6E
74 65 6E 74 73 BF 63 76 61 6C 78 18 6C 6F 74 73
20 6F 66 20 68 69 62 65 72 6E 61 74 69 6F 6E 20
64 61 74 61 FF 64 6E 61 6D 65 BF 63 76 61 6C 6C
68 69 62 65 72 66 69 6C 2E 73 79 73 FF FF BF 62
40 54 61 31 62 40 69 07 62 40 74 64 46 69 6C 65
68 63 6F 6E 74 65 6E 74 73 BF 63 76 61 6C 76 6C
6F 74 73 20 6F 66 20 70 61 67 65 20 66 69 6C 65
20 64 61 74 61 FF 64 6E 61 6D 65 BF 63 76 61 6C
6C 70 61 67 65 66 69 6C 65 2E 73 79 73 FF FF BF
62 40 54 61 31 62 40 69 08 62 40 74 64 46 69 6C
65 68 63 6F 6E 74 65 6E 74 73 BF 63 76 61 6C 71
6C 6F 74 73 20 6F 66 20 73 77 61 70 20 64 61 74
61 FF 64 6E 61 6D 65 BF 63 76 61 6C 6C 73 77 61
70 66 69 6C 65 2E 73 79 73 FF FF BF 62 40 54 61
31 62 40 69 09 62 40 74 65 4D 6F 75 6E 74 66 74
61 72 67 65 74 BF 62 40 54 61 24 62 40 6C 0B FF
64 6E 61 6D 65 BF 63 76 61 6C 68 53 6F 6D 65 55
73 65 72 FF FF FF 64 6E 61 6D 65 BF 63 76 61
6C 60 FF FF FF BF 62 40 54 61 31 62 40 69 0A 62
40 74 65 44 72 69 76 65 66 6C 65 74 74 65 72 BF
```

(continues on next page)

(continued from previous page)

```

63 76 61 6C 18 44 FF 68 72 6F 6F 74 5F 64 69 72
BF 62 40 54 61 31 62 40 69 0B 62 40 74 69 44 69
72 65 63 74 6F 72 79 67 65 6E 74 72 69 65 73 BF
62 40 54 61 2A 62 40 64 9F BF 62 40 54 61 31 62
40 69 0C 62 40 74 65 4D 6F 75 6E 74 66 74 61 72
67 65 74 BF 62 40 54 61 24 62 40 6C 02 FF 64 6E
61 6D 65 BF 63 76 61 6C 6F 65 76 69 6C 20 6C 69
6E 6B 20 74 6F 20 43 3A FF FF FF FF 64 6E 61 6D
65 BF 63 76 61 6C 60 FF FF FF FF FF

```

You can pull that through <http://cbor.me> and <https://jsonformatter.org> to inspect the output, if you like.

We can deserialize it into a complete copy of the tree.

```

auto system2 = tree::base::deserialize<directory::System>(cbor);
system2->dump();

```

↓

```

System(
  drives: [
    Drive(
      letter: C
      root_dir: <
        Directory(
          entries: [
            Directory(
              entries: []
              name: Program Files
            )
            Directory(
              entries: []
              name: Windows
            )
            Directory(
              entries: []
              name: Users
            )
            File(
              contents: lots of hibernation data
              name: hiberfil.sys
            )
            File(
              contents: lots of page file data
              name: pagefile.sys
            )
            File(
              contents: lots of swap data
              name: swapfile.sys
            )
            Mount(
              target --> <
                Directory(
                  entries: [
                    Mount(
                      target --> <
                        ...

```

(continues on next page)

(continued from previous page)

```
        >
        name: evil link to C:
      )
    ]
    name:
  )
  >
  name: SomeUser
)
]
name:
)
>
)
Drive(
  letter: D
  root_dir: <
  Directory(
    entries: [
      Mount(
        target --> <
        Directory(
          entries: [
            Directory(
              entries: []
              name: Program Files
            )
            Directory(
              entries: []
              name: Windows
            )
            Directory(
              entries: []
              name: Users
            )
            File(
              contents: lots of hibernation data
              name: hiberfil.sys
            )
            File(
              contents: lots of page file data
              name: pagefile.sys
            )
            File(
              contents: lots of swap data
              name: swapfile.sys
            )
            Mount(
              target --> <
              ...
            >
            name: SomeUser
          )
        ]
        name:
      )
    ]
  >
)
>
```

(continues on next page)

(continued from previous page)

```

        name: evil link to C:
    )
]
name:
)
>
)
]
)

```

Note that equality for two link edges is satisfied only if they point to the exact same node. That's not the case for the links in our two entirely separate trees, so the two trees register as unequal.

```
ASSERT(!system2.equals(system));
```

To be sure no data was lost, we'll have to check the CBOR and debug dumps instead.

```

ASSERT(tree::base::serialize(system) == tree::base::serialize(system2));
std::ostream ss1{};
system->dump(ss1);
std::ostream ss2{};
system2->dump(ss2);
ASSERT(ss1.str() == ss2.str());

```

2.1 Python usage

Let us now turn our attention to tree-gen's Python support. As you may remember from the introduction, tree-gen does not provide a direct interface between Python and C++ using a tool like SWIG; rather, it provides conversion between a CBOR tree representation and the in-memory representations of both languages. It then becomes trivial for users of the tree-gen structure to provide an API to users that converts from one to the other, since only a single string of bytes has to be transferred.

Recall that we wrote the CBOR representation of the tree we constructed in the C++ example to a file. Let's try loading this file with Python.

```

from directory import *

with open(os.path.join(TEST_DIR, 'tree.cbor'), 'rb') as f:
    tree = System.deserialize(f.read())

tree.check_well_formed()
print(tree)

```

↓

```

System(
  drives: [
    Drive(
      letter: C
      root_dir: <
        Directory(
          entries: [
            Directory(
              entries: -

```

(continues on next page)

(continued from previous page)

```

        name: Program Files
    )
    Directory(
        entries: -
        name: Windows
    )
    Directory(
        entries: -
        name: Users
    )
    File(
        contents: lots of hibernation data
        name: hiberfil.sys
    )
    File(
        contents: lots of page file data
        name: pagefile.sys
    )
    File(
        contents: lots of swap data
        name: swapfile.sys
    )
    Mount(
        target --> <
            Directory(
                entries: [
                    Mount(
                        target --> <
                            ...
                        >
                        name: evil link to C:
                    )
                ]
                name:
            )
        >
        name: SomeUser
    )
]
name:
)
>
)
Drive(
    letter: D
    root_dir: <
        Directory(
            entries: [
                Mount(
                    target --> <
                        Directory(
                            entries: [
                                Directory(
                                    entries: -
                                    name: Program Files
                                )
                            ]
                        )
                    )
                ]
            )
        )
    )

```

(continues on next page)


```
Traceback (most recent call last):
  File "main.py", line 48, in <module>
    tree.color_of_a_banana = 'yellow'
AttributeError: 'System' object has no attribute 'color_of_a_banana'
```

... or assign values of incorrect types to ones that do:

```
tree.drives[0] = 'not-a-drive'
```

↓

```
Traceback (most recent call last):
  File "main.py", line 55, in <module>
    tree.drives[0] = 'not-a-drive'
  File "/home/docs/checkouts/readthedocs.org/user_builds/tree-gen/checkouts/latest/
↳cbuild/examples/directory/directory.py", line 490, in __setitem__
    .format(val, idx, self._T)
TypeError: object 'not-a-drive' is not an instance of 0
```

Note that the setters will however try to “typecast” objects they don’t know about (to be more specific, anything that isn’t an instance of Node). So if we try to assign a string directly to `tree.drives`, you get something you might not have expected:

```
tree.drives = 'not-a-set-of-drives'
```

↓

```
Traceback (most recent call last):
  File "main.py", line 65, in <module>
    tree.drives = 'not-a-set-of-drives'
  File "/home/docs/checkouts/readthedocs.org/user_builds/tree-gen/checkouts/latest/
↳cbuild/examples/directory/directory.py", line 1576, in drives
    val = MultiDrive(val)
  File "/home/docs/checkouts/readthedocs.org/user_builds/tree-gen/checkouts/latest/
↳cbuild/examples/directory/directory.py", line 472, in __init__
    .format(val, idx, self._T)
TypeError: object 'n' at index 0 is not an instance of <class 'directory.Drive'>
```

The setter for the `drives` property is trying to cast the string to a `MultiDrive`, which interprets its string input as any other iterable, and ultimately tries to cast the first letter to a `Drive`. There’s not much that can be done about this; it’s just how duck typing works.

Manipulation of trees in Python works just as you might expect it to in Python:

```
tree.drives.append(Drive('E'))
tree.drives[-1].root_dir = Directory([
    File('test', 'test-contents'),
    File('test2', 'test2-contents'),
])
del tree.drives[1]
assert not tree.is_well_formed()
tree.drives[0].root_dir.entries[-1].target = tree.drives[1].root_dir
tree.check_well_formed()
print(tree)
```

↓

```
System(  
  drives: [  
    Drive(  
      letter: C  
      root_dir: <  
        Directory(  
          entries: [  
            Directory(  
              entries: -  
              name: Program Files  
            )  
            Directory(  
              entries: -  
              name: Windows  
            )  
            Directory(  
              entries: -  
              name: Users  
            )  
            File(  
              contents: lots of hibernation data  
              name: hiberfil.sys  
            )  
            File(  
              contents: lots of page file data  
              name: pagefile.sys  
            )  
            File(  
              contents: lots of swap data  
              name: swapfile.sys  
            )  
            Mount(  
              target --> <  
                Directory(  
                  entries: [  
                    File(  
                      contents: test  
                      name: test-contents  
                    )  
                    File(  
                      contents: test2  
                      name: test2-contents  
                    )  
                  ]  
                  name:  
                )  
              >  
              name: SomeUser  
            )  
          ]  
          name:  
        )  
      >  
    )  
    Drive(  
      letter: E  
      root_dir: <
```

(continues on next page)

(continued from previous page)

```

    Directory(
        entries: [
            File(
                contents: test
                name: test-contents
            )
            File(
                contents: test2
                name: test2-contents
            )
        ]
        name:
    )
>
)
]
)

```

Once the Python code is done manipulating the structure, it can be serialized again.

```

cbor = tree.serialize()
assert cbor == System.deserialize(cbor).serialize()
count = 0
for c in cbor:
    if count == 16:
        print()
        count = 0
    elif count > 0 and count % 4 == 0:
        print(' ', end='')
    print('%02X ' % c, end='')
    count += 1
print()

```

↓

```

A3 62 40 69 00 62 40 74 66 53 79 73 74 65 6D 66
64 72 69 76 65 73 A2 62 40 54 61 2B 62 40 64 82
A5 62 40 54 61 31 62 40 69 01 62 40 74 65 44 72
69 76 65 66 6C 65 74 74 65 72 A1 63 76 61 6C 18
43 68 72 6F 6F 74 5F 64 69 72 A5 62 40 54 61 31
62 40 69 02 62 40 74 69 44 69 72 65 63 74 6F 72
79 67 65 6E 74 72 69 65 73 A2 62 40 54 61 2A 62
40 64 87 A5 62 40 54 61 31 62 40 69 03 62 40 74
69 44 69 72 65 63 74 6F 72 79 67 65 6E 74 72 69
65 73 A2 62 40 54 61 2A 62 40 64 80 64 6E 61 6D
65 A1 63 76 61 6C 6D 50 72 6F 67 72 61 6D 20 46
69 6C 65 73 A5 62 40 54 61 31 62 40 69 04 62 40
74 69 44 69 72 65 63 74 6F 72 79 67 65 6E 74 72
69 65 73 A2 62 40 54 61 2A 62 40 64 80 64 6E 61
6D 65 A1 63 76 61 6C 67 57 69 6E 64 6F 77 73 A5
62 40 54 61 31 62 40 69 05 62 40 74 69 44 69 72
65 63 74 6F 72 79 67 65 6E 74 72 69 65 73 A2 62
40 54 61 2A 62 40 64 80 64 6E 61 6D 65 A1 63 76
61 6C 65 55 73 65 72 73 A5 62 40 54 61 31 62 40
69 06 62 40 74 64 46 69 6C 65 68 63 6F 6E 74 65
6E 74 73 A1 63 76 61 6C 78 18 6C 6F 74 73 20 6F
66 20 68 69 62 65 72 6E 61 74 69 6F 6E 20 64 61

```

(continues on next page)

(continued from previous page)

```

74 61 64 6E 61 6D 65 A1 63 76 61 6C 6C 68 69 62
65 72 66 69 6C 2E 73 79 73 A5 62 40 54 61 31 62
40 69 07 62 40 74 64 46 69 6C 65 68 63 6F 6E 74
65 6E 74 73 A1 63 76 61 6C 76 6C 6F 74 73 20 6F
66 20 70 61 67 65 20 66 69 6C 65 20 64 61 74 61
64 6E 61 6D 65 A1 63 76 61 6C 6C 70 61 67 65 66
69 6C 65 2E 73 79 73 A5 62 40 54 61 31 62 40 69
08 62 40 74 64 46 69 6C 65 68 63 6F 6E 74 65 6E
74 73 A1 63 76 61 6C 71 6C 6F 74 73 20 6F 66 20
73 77 61 70 20 64 61 74 61 64 6E 61 6D 65 A1 63
76 61 6C 6C 73 77 61 70 66 69 6C 65 2E 73 79 73
A5 62 40 54 61 31 62 40 69 09 62 40 74 65 4D 6F
75 6E 74 64 6E 61 6D 65 A1 63 76 61 6C 68 53 6F
6D 65 55 73 65 72 66 74 61 72 67 65 74 A2 62 40
54 61 24 62 40 6C 0B 64 6E 61 6D 65 A1 63 76 61
6C 60 A5 62 40 54 61 31 62 40 69 0A 62 40 74 65
44 72 69 76 65 66 6C 65 74 74 65 72 A1 63 76 61
6C 18 45 68 72 6F 6F 74 5F 64 69 72 A5 62 40 54
61 31 62 40 69 0B 62 40 74 69 44 69 72 65 63 74
6F 72 79 67 65 6E 74 72 69 65 73 A2 62 40 54 61
2A 62 40 64 82 A5 62 40 54 61 31 62 40 69 0C 62
40 74 64 46 69 6C 65 68 63 6F 6E 74 65 6E 74 73
A1 63 76 61 6C 64 74 65 73 74 64 6E 61 6D 65 A1
63 76 61 6C 6D 74 65 73 74 2D 63 6F 6E 74 65 6E
74 73 A5 62 40 54 61 31 62 40 69 0D 62 40 74 64
46 69 6C 65 68 63 6F 6E 74 65 6E 74 73 A1 63 76
61 6C 65 74 65 73 74 32 64 6E 61 6D 65 A1 63 76
61 6C 6E 74 65 73 74 32 2D 63 6F 6E 74 65 6E 74
73 64 6E 61 6D 65 A1 63 76 61 6C 60

```

2.2 Complete file listings

2.2.1 directory.tree

```

1 // Attach \file docstrings to the generated files for Doxygen.
2 # Implementation for classes representing a Windows directory tree.
3 source
4
5 # Header for classes representing a Windows directory tree.
6 header
7
8 // Include tree base classes.
9 include "tree-base.hpp"
10 tree_namespace tree::base
11
12 // Include primitive types.
13 include "primitives.hpp"
14 import primitives
15
16 // Initialization function to use to construct default values for the tree base
17 // classes and primitives.
18 initialize_function primitives::initialize
19 serdes_functions primitives::serialize primitives::deserialize

```

(continues on next page)

(continued from previous page)

```
20
21 // Set the namespace for the generated classes and attach a docstring.
22 # Namespace for classes representing a Windows directory tree.
23 namespace directory
24
25 # Root node, containing the drives and associated directory trees.
26 system {
27
28     # The drives available on the system. There must be at least one.
29     drives: Many<drive>;
30
31 }
32
33 # Represents a drive.
34 drive {
35
36     # The drive letter used to identify it.
37     letter: primitives::Letter;
38
39     # Root directory.
40     root_dir: One<directory>;
41
42 }
43
44 # Represents a directory entry.
45 entry {
46
47     # Name of the directory entry.
48     name: primitives::String;
49
50     # Represents a regular file.
51     file {
52
53         # The file contents.
54         contents: primitives::String;
55
56     }
57
58     # Represents a (sub)directory.
59     directory {
60
61         # The directory contents. Note that directories can be empty.
62         entries: Any<entry>;
63
64     }
65
66     # Represents a link to another directory.
67     mount {
68
69         # The directory linked to.
70         target: Link<directory>;
71
72     }
73
74 }
```

2.2.2 primitives.hpp

```
1  /** \file
2   * Defines primitives used in the generated directory tree structure.
3   */
4
5  #pragma once
6
7  #include <string>
8
9  /**
10   * Namespace with primitives used in the generated directory tree structure.
11   */
12  namespace primitives {
13
14   /**
15    * Letter primitive, used to represent drive letters.
16    */
17   using Letter = char;
18
19   /**
20    * Strings, used to represent filenames and file contents.
21    */
22   using String = std::string;
23
24   /**
25    * Initialization function. This must be specialized for any types used as
26    * primitives in a tree that are actual C primitives (int, char, bool, etc),
27    * as these are not initialized by the T() construct.
28    */
29   template <class T>
30   T initialize() { return T(); };
31
32   /**
33    * Declare the default initializer for drive letters. It's declared inline
34    * to avoid having to make a cpp file just for this.
35    */
36   template <>
37   inline Letter initialize<Letter>() {
38       return 'A';
39   }
40
41   /**
42    * Serialization function. This must be specialized for any types used as
43    * primitives in a tree. The default implementation doesn't do anything.
44    */
45   template <typename T>
46   void serialize(const T &obj, tree::cbor::MapWriter &map) {
47   }
48
49   /**
50    * Serialization function for Letter.
51    */
52   template <>
53   inline void serialize<Letter>(const Letter &obj, tree::cbor::MapWriter &map) {
54       map.append_int("val", obj);
55   }
```

(continues on next page)

(continued from previous page)

```

56
57 /**
58  * Serialization function for String.
59  */
60 template <>
61 inline void serialize<String>(const String &obj, tree::cbor::MapWriter &map) {
62     map.append_string("val", obj);
63 }
64
65 /**
66  * Deserialization function. This must be specialized for any types used as
67  * primitives in a tree. The default implementation doesn't do anything.
68  */
69 template <typename T>
70 T deserialize(const tree::cbor::MapReader &map) {
71     return initialize<T>();
72 }
73
74 /**
75  * Deserialization function for Letter.
76  */
77 template <>
78 inline Letter deserialize<Letter>(const tree::cbor::MapReader &map) {
79     return map.at("val").as_int();
80 }
81
82 /**
83  * Deserialization function for String.
84  */
85 template <>
86 inline String deserialize<String>(const tree::cbor::MapReader &map) {
87     return map.at("val").as_string();
88 }
89
90 } // namespace primitives

```

2.2.3 primitives.py

```

1 class Letter(str):
2     """Letter primitive, used to represent drive letters."""
3     def __new__(cls, s='A'):
4         return str.__new__(cls, s)
5
6
7 class String(str):
8     """Strings, used to represent filenames and file contents."""
9     pass
10
11
12 def serialize(typ, val):
13     """Serialization function."""
14
15     # Serialization formats of primitives.
16     if typ is Letter:
17         return {'val': ord(val)}

```

(continues on next page)

(continued from previous page)

```
18     if typ is String:
19         return {'val': val}
20
21     # Serialization formats of annotations.
22     if isinstance(typ, str):
23         return val
24
25     # Some unknown type.
26     raise TypeError('no known serialization of type %r, value %r' % (typ, val))
27
28
29 def deserialize(typ, val):
30     """Deserialization function."""
31
32     # Serialization formats of primitives.
33     if typ is Letter:
34         return Letter(chr(val['val']))
35     if typ is String:
36         return String(val['val'])
37
38     # Serialization formats of annotations.
39     if isinstance(typ, str):
40         return val
41
42     # Some unknown type.
43     raise TypeError('no known deserialization for type %r, value %r' % (typ, val))
```

2.2.4 CMakeLists.txt

```
1  cmake_minimum_required(VERSION 3.12 FATAL_ERROR)
2
3  project(directory-example CXX)
4
5  # Add the tree-gen repository root directory. Normally, your project would be
6  # located outside of tree-gen's root, which means you can omit the second
7  # argument.
8  add_subdirectory(..../.. tree-gen)
9
10 # Generates the files for the directory tree.
11 generate_tree_py(
12     "${CMAKE_CURRENT_SOURCE_DIR}/directory.tree"
13     "${CMAKE_CURRENT_BINARY_DIR}/directory.hpp"
14     "${CMAKE_CURRENT_BINARY_DIR}/directory.cpp"
15     "${CMAKE_CURRENT_BINARY_DIR}/directory.py"
16 )
17
18 add_executable(
19     directory-example
20     "${CMAKE_CURRENT_BINARY_DIR}/directory.cpp"
21     "${CMAKE_CURRENT_SOURCE_DIR}/main.cpp"
22 )
23
24 target_include_directories(
25     directory-example
26     # This directory for primitives.hpp:
```

(continues on next page)

(continued from previous page)

```

27     PRIVATE "${CMAKE_CURRENT_SOURCE_DIR}"
28     # Binary directory for directory.hpp:
29     PRIVATE "${CMAKE_CURRENT_BINARY_DIR}"
30 )
31
32 target_link_libraries(directory-example tree-lib)
33
34 # The following lines only serve to register the example as a test, so you can
35 # run it using `make test` or `ctest` as well. You don't need them as such in
36 # your own project.
37 enable_testing()
38 add_test(
39     NAME directory-example
40     COMMAND directory-example
41     WORKING_DIRECTORY ${CMAKE_CURRENT_BINARY_DIR}
42 )
43
44 # Only add the Python test if CMake is new enough for us to not have to bother
45 # with FindPythonInterp.
46 if(NOT ${CMAKE_VERSION} VERSION_LESS "3.12")
47     find_package(Python3 COMPONENTS Interpreter)
48     if(${Python3_FOUND})
49         add_test(
50             NAME directory-example-py
51             COMMAND ${Python3_EXECUTABLE} main.py ${CMAKE_CURRENT_BINARY_DIR}
52             WORKING_DIRECTORY ${CMAKE_CURRENT_SOURCE_DIR}
53         )
54     endif()
55 endif()

```

2.2.5 main.cpp

```

1  #include <iostream>
2  #include <cstdio>
3  #include <stdexcept>
4  #include "../utils.hpp"
5
6  // Include the generated file.
7  #include "directory.hpp"
8
9  // Note: the // comment contents of main(), together with the MARKER lines and
10 // the output of the program, are used to automatically turn this into a
11 // restructured-text page for ReadTheDocs.
12
13 int main() {
14
15     // *****
16     // Directory tree example
17     // *****
18     //
19     // This example illustrates the tree system with a Windows-like directory
20     // tree structure. The ``System`` root node consists of one or more drives,
21     // each of which has a drive letter and a root directory with named files
22     // and subdirectories. To make things a little more interesting, a
23     // symlink-like "mount" is added as a file type, which links to another

```

(continues on next page)

```
24 // directory.
25 //
26 // Using this tree, this example should teach you the basics of using
27 // tree-gen trees. The tutorial runs you through the C++ code of
28 // ``main.cpp`` and Python code of ``main.py``, but be sure to check out
29 // ``directory.tree``, ``primitives.hpp``, ``primitives.hpp``, and (if you
30 // want to reproduce it in your own project) ``CMakeLists.txt`` copied to
31 // the bottom of this page as well. You will also find the complete
32 // ``main.cpp`` and ``main.py`` there.
33 //
34 // Let's start by making the root node using ``tree::base::make()``. This
35 // works somewhat like ``std::make_shared()`` , but instead of returning a
36 // ``shared_ptr`` it returns a ``One`` edge. This might come off as a bit
37 // odd, considering trees in graph theory start with a node rather than an
38 // edge. This is just a choice, though; a side effect of how the internal
39 // ``shared_ptr``'s work and how Link/OptLink edges work (if you'd store the
40 // tree as the root structure directly, the root node wouldn't always be
41 // stored in the same place on the heap, breaking link nodes).
42 auto system = tree::base::make<directory::System>();
43 MARKER
44
45 // At all times, you can use the ``dump()`` method on a node to see what it
46 // looks like for debugging purposes. It looks like this:
47 system->dump();
48 MARKER
49
50 // While this system node exists as a tree in memory and tree-gen seems
51 // happy, our system tree is at this time not "well-formed". A tree node (or
52 // an edge to one) is only considered well-formed if all of the following
53 // are true:
54 //
55 // - All ``One`` edges in the tree connect to exactly one node.
56 // - All ``Many`` edges in the tree connect to at least one node.
57 // - All ``Link`` and non-empty ``OptLink`` nodes refer to a node reachable
58 //   from the root node.
59 // - Each node in the tree is only used by a non-link node once.
60 //
61 // Currently, the second requirement is not met.
62 ASSERT(!system.is_well_formed());
63 MARKER
64
65 // You can get slightly more information using ``check_well_formed()``;
66 // instead of returning a boolean, it will throw a ``NotWellFormed``
67 // exception with an error message.
68 ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
69 MARKER
70
71 // Note that the name of the type is
72 // "mangled" <https://en.wikipedia.org/wiki/Name\_mangling#C++>`. The exact
73 // output will vary from compiler to compiler, or even from project to
74 // project. But hopefully it'll be readable enough to make sense of in
75 // general.
76
77 // To fix that, let's add a default drive node to the tree. This should get
78 // drive letter ``A``, because primitives::initialize() is specialized to
79 // return that for ``Letter``'s (see primitives.hpp).
80 system->drives.add(tree::base::make<directory::Drive>());
```

(continues on next page)

(continued from previous page)

```

81  MARKER
82
83  // ``Drive`` has a ``One`` edge that is no empty, though, so the tree still
84  // isn't well-formed. Let's add one of those as well.
85  system->drives[0]->root_dir = tree::base::make<directory::Directory>();
86  MARKER
87
88  // Now we have a well-formed tree. Let's have a look:
89  system.check_well_formed();
90  system->dump();
91  MARKER
92
93  // We can just change the drive letter by assignment, as you would expect.
94  system->drives[0]->letter = 'C';
95  MARKER
96
97  // Before we add files and directories, let's make a shorthand variable for
98  // the root directory. Because root_dir is an edge to another node rather
99  // than the node itself, and thus acts like a pointer or reference to it,
100 // we can just copy it into a variable and modify the variable to update the
101 // tree. Note that you'd normally just write "auto" for the type for
102 // brevity; the type is shown here to make explicit what it turns into.
103 directory::One<directory::Directory> root = system->drives[0]->root_dir;
104 MARKER
105
106 // Let's make a "Program Files" subdirectory and add it to the root.
107 auto programs = tree::base::make<directory::Directory>(
108     tree::base::Any<directory::Entry>{},
109     "Program Files");
110 root->entries.add(programs);
111 system.check_well_formed();
112 MARKER
113
114 // That's quite verbose. But in most cases it can be written way shorter.
115 // Here's the same with the less versatile but also less verbose emplace()
116 // method (which avoids the tree::make() call, but doesn't allow an
117 // insertion index to be specified) and with a "using namespace" for the
118 // tree. emplace() can also be chained, allowing multiple files and
119 // directories to be added at once in this case.
120 {
121     using namespace directory;
122
123     root->entries.emplace<Directory>(Any<Entry>{}, "Windows")
124         .emplace<Directory>(Any<Entry>{}, "Users")
125         .emplace<File>("lots of hibernation data", "hiberfil.sys")
126         .emplace<File>("lots of page file data", "pagefile.sys")
127         .emplace<File>("lots of swap data", "swapfile.sys");
128 }
129 system.check_well_formed();
130 MARKER
131
132 // In order to look for a file in a directory, you'll have to make your own
133 // function to iterate over them. After all, tree-gen doesn't know that the
134 // name field is a key; it has no concept of a key-value store. This is
135 // simple enough to make, but to prevent this example from getting out of
136 // hand we'll just use indices for now.
137

```

(continues on next page)

(continued from previous page)

```

138 // Let's try to read the "lots of data" string from pagefile.sys.
139 ASSERT(root->entries[4]->name == "pagefile.sys");
140 // ASSERT(root->entries[4]->contents == "lots of page file data");
141 // '-> No member named 'contents' in 'directory::Entry'
142 MARKER
143
144 // Hm, that didn't work so well, because C++ doesn't know that entry 4
145 // happens to be a file rather than a directory or a mount. We have to tell
146 // it to cast to a file first (which throws a std::bad_cast if it's not
147 // actually a file). The easiest way to do that is like this:
148 ASSERT(root->entries[4]->as_file()->contents == "lots of page file data");
149 MARKER
150
151 // No verbose casts required; tree-gen will make member functions for all
152 // the possible subtypes.
153
154 // While it's possible to put the same node in a tree twice (without using
155 // a link), this is not allowed. This isn't checked until a well-formedness
156 // check is performed, however (and in fact can't be without having access
157 // to the root node).
158 root->entries.add(root->entries[0]);
159 ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
160 MARKER
161
162 // Note that we can index nodes Python-style with negative integers for
163 // add() and remove(), so remove(-1) will remove the broken node we just
164 // added. Note that the -1 is not actually necessary, though, as it is the
165 // default.
166 root->entries.remove(-1);
167 system.check_well_formed();
168 MARKER
169
170 // We *can*, of course, add copies of nodes. That's what copy (shallow) and
171 // clone (deep) are for. Usually you'll want a deep copy, but in this case
172 // shallow is fine, because a File has no child nodes. Note that, even for
173 // a deep copy, links are not modified; it is intended to copy a subtree to
174 // another part of the same tree. To make a complete copy of a tree that
175 // maintains link semantics (i.e. does not link back to the original tree)
176 // the best way would be to go through serialize/deserialize.
177 root->entries.add(root->entries[0]->copy());
178 system.check_well_formed();
179 MARKER
180
181 // Note that the generated classes don't care that there are now two
182 // directories named Program Files in the root. As far as they're concerned,
183 // they're two different directories with the same name. Let's remove it
184 // again, though.
185 root->entries.remove();
186 MARKER
187
188 // Something we haven't looked at yet are links. Links are edges in the tree
189 // that, well, turn it into something that isn't strictly a tree anymore.
190 // While One/Maybe/Any/Many require that nodes are unique, Link/OptLink
191 // require that they are *not* unique, and are present elsewhere in the
192 // tree. Let's make a new directory, and mount it in the Users directory.
193 auto user_dir = tree::base::make<directory::Directory>(
194     tree::base::Any<directory::Entry>{}),

```

(continues on next page)

(continued from previous page)

```

195     "");
196     root->entries.emplace<directory::Mount>(user_dir, "SomeUser");
197     MARKER
198
199     // Note that user_dir is not yet part of the tree. emplace works simply
200     // because it doesn't check whether the directory is in the tree yet. But
201     // the tree is no longer well-formed now.
202     ASSERT_RAISES(tree::base::NotWellFormed, system.check_well_formed());
203     MARKER
204
205     // To make it work again, we can add it as a root directory to a second
206     // drive.
207     system->drives.emplace<directory::Drive>('D', user_dir);
208     system.check_well_formed();
209     MARKER
210
211     // A good way to confuse a filesystem is to make loops. tree-gen is okay
212     // with this, though.
213     system->drives[1]->root_dir->entries.emplace<directory::Mount>(root, "evil link_
↳to C:");
214     system.check_well_formed();
215     MARKER
216
217     // The only place where it matters is in the dump function, which only goes
218     // one level deep. After that, it'll just print an ellipsis.
219     system->dump();
220     MARKER
221
222     // Now that we have a nice tree, let's try the serialization and
223     // deserialization functionality. Serializing is easy:
224     std::string cbor = tree::base::serialize(system);
225     MARKER
226
227     // Let's write it to a file; we'll load this in Python later.
228     {
229         std::ofstream cbor_output;
230         cbor_output.open("tree.cbor", std::ios::out | std::ios::trunc |
↳std::ios::binary);
231         cbor_output << cbor;
232     }
233     MARKER
234
235     // Let's also have a look at a hexdump of that.
236     int count = 0;
237     for (char c : cbor) {
238         if (count == 16) {
239             std::printf("\n");
240             count = 0;
241         } else if (count > 0 && count % 4 == 0) {
242             std::printf(" ");
243         }
244         std::printf("%02X ", (uint8_t)c);
245         count++;
246     }
247     std::printf("\n");
248     MARKER
249

```

(continues on next page)

(continued from previous page)

```

250 // You can pull that through http://cbor.me and https://jsonformatter.org
251 // to inspect the output, if you like.
252
253 // We can deserialize it into a complete copy of the tree.
254 auto system2 = tree::base::deserialize<directory::System>(cbor);
255 system2->dump();
256 MARKER
257
258 // Note that equality for two link edges is satisfied only if they point to
259 // the exact same node. That's not the case for the links in our two
260 // entirely separate trees, so the two trees register as unequal.
261 ASSERT(!system2.equals(system));
262 MARKER
263
264 // To be sure no data was lost, we'll have to check the CBOR and debug dumps
265 // instead.
266 ASSERT(tree::base::serialize(system) == tree::base::serialize(system2));
267 std::ostringstream ss1{};
268 system->dump(ss1);
269 std::ostringstream ss2{};
270 system2->dump(ss2);
271 ASSERT(ss1.str() == ss2.str());
272 MARKER
273
274 return 0;
275 }

```

2.2.6 main.py

```

1 import sys, os, traceback
2 TEST_DIR = os.path.realpath(sys.argv[1])
3 sys.path.append(TEST_DIR)
4
5
6 def marker():
7     print('###MARKER###')
8
9
10 # Note: the # | comment contents of this file, together with the marker() lines
11 # and the output of the program, are used to automatically turn this into a
12 # restructured-text page for ReadTheDocs. The output is appended to the C++
13 # output.
14
15 # | Python usage
16 # | =====
17
18 # | Let us now turn our attention to tree-gen's Python support. As you may
19 # | remember from the introduction, tree-gen does not provide a direct interface
20 # | between Python and C++ using a tool like SWIG; rather, it provides
21 # | conversion between a CBOR tree representation and the in-memory
22 # | representations of both languages. It then becomes trivial for users of the
23 # | tree-gen structure to provide an API to users that converts from one to the
24 # | other, since only a single string of bytes has to be transferred.
25
26 # | Recall that we wrote the CBOR representation of the tree we constructed in

```

(continues on next page)

(continued from previous page)

```
27 # | the C++ example to a file. Let's try loading this file with Python.
28 from directory import *
29
30 with open(os.path.join(TEST_DIR, 'tree.cbor'), 'rb') as f:
31     tree = System.deserialize(f.read())
32
33 tree.check_well_formed()
34 print(tree)
35 marker()
36
37 # | And there we go; the same structure in pure Python.
38
39 # | As you might expect from a pure-Python structure, field access is much less
40 # | verbose than it is in C++.
41 print(tree.drives[0].letter)
42 print(tree.drives[0].root_dir.entries[0].name)
43 marker()
44
45 # | Perhaps contrary to Python intuition however, the structure provides
46 # | type-safety. You can't assign values to fields that don't exist...
47 try:
48     tree.color_of_a_banana = 'yellow'
49 except AttributeError:
50     traceback.print_exc(file=sys.stdout)
51 marker()
52
53 # | ... or assign values of incorrect types to ones that do:
54 try:
55     tree.drives[0] = 'not-a-drive'
56 except TypeError:
57     traceback.print_exc(file=sys.stdout)
58 marker()
59
60 # | Note that the setters will however try to "typecast" objects they don't know
61 # | about (to be more specific, anything that isn't an instance of Node). So if
62 # | we try to assign a string directly to tree.drives, you get something you
63 # | might not have expected:
64 try:
65     tree.drives = 'not-a-set-of-drives'
66 except TypeError:
67     traceback.print_exc(file=sys.stdout)
68 marker()
69
70 # | The setter for the drives property is trying to cast the string to a
71 # | MultiDrive, which interprets its string input as any other iterable, and
72 # | ultimately tries to cast the first letter to a Drive. There's not much that
73 # | can be done about this; it's just how duck typing works.
74
75 # | Manipulation of trees in Python works just as you might expect it to in
76 # | Python:
77 tree.drives.append(Drive('E'))
78 tree.drives[-1].root_dir = Directory([
79     File('test', 'test-contents'),
80     File('test2', 'test2-contents'),
81 ])
82 del tree.drives[1]
83 assert not tree.is_well_formed()
```

(continues on next page)

(continued from previous page)

```
84 tree.drives[0].root_dir.entries[-1].target = tree.drives[1].root_dir
85 tree.check_well_formed()
86 print(tree)
87 marker()
88
89 # | Once the Python code is done manipulating the structure, it can be
90 # | serialized again.
91 cbor = tree.serialize()
92 assert cbor == System.deserialize(cbor).serialize()
93 count = 0
94 for c in cbor:
95     if count == 16:
96         print()
97         count = 0
98     elif count > 0 and count % 4 == 0:
99         print(' ', end='')
100         print('%02X ' % c, end='')
101         count += 1
102 print()
103 marker()
```

Interpreter example

This example discusses some of the more advanced features of `tree-gen`, using an interpreter for a very simple language as an example. This is, however, not a completely integrated or even functional interpreter; we just go over some of the concepts and things you may run into when you would make such an interpreter.

3.1 External nodes & multiple trees

When trees get large and complicated, you may want to split a tree up into multiple files. This is supported within `tree-gen`, as long as there are no circular references to nodes between the files. In this example, we've split up the expression-like nodes into a tree separate from the statement-like nodes. The former, `value.tree`, has nothing special going on in this regard, while the latter, `program.tree`, effectively includes the former. The files are listed at the bottom of this page for your convenience.

`tree-gen` itself is unaware of any links between files. For `program.tree`, it simply generates code assuming that the types referred to through the `external` edges exist. For this to be true, you have to tell `tree-gen` to include the external tree's generated header file at the top of its generated header file using the `include` directive. Note also that, as `tree-gen` is unaware of the external tree, you'll have to use the full C++ `TitleCase` type name, rather than the `snake_case` name for normal edges.

3.2 Integration with Flex/Bison (or other parsers)

When making an interpreter (or parser of any kind) with `tree-gen`, you'll probably want to use a tokenizer/parser generator such as Flex/Bison. With Bison at least, you'll quickly run into a fundamental issue when trying to store `tree-gen` nodes in the generated `yyval` union: union in C++ can only consist of types with trivial constructors, and `tree-gen` nodes are not among those. The author is unfamiliar with other frameworks like it, but this will probably apply for any C-style parser generator.

The solution is to use raw pointers to the generated node types in this enumeration. For instance, you may end up with the following (which would be generated by Bison):

```

union node_type {
    program::Program *prog;
    program::Assignment *asgn;
    value::Literal *lit;
    value::Reference *ref;
    // ...
};

```

To create a new node, instead of using `tree::base::make`, you use the `new` keyword directly.

```

node_type n1, n2;
n1.lit = new value::Literal(2);
n2.ref = new value::Reference("x");

```

The generated constructors won't help you much for non-leaf nodes, because they expect edge wrappers (`One`, `Maybe`, etc.) around the parameters, so you'll typically want to use the constructor without arguments. Instead, to assign child nodes, you can use `Maybe::set_raw()`, `One::set_raw()`, `Any::add_raw()`, or `Many::add_raw()`. These functions take ownership of a new-allocated raw pointer to a node - exactly what we need here. Note that because they take ownership, you don't have to (and shouldn't!) delete manually.

```

node_type n3;
n3.asgn = new program::Assignment();
n3.asgn->rhs.set_raw(n1.lit);
n3.asgn->lhs.set_raw(n2.ref);

node_type n4;
n4.prog = new program::Program();
n4.prog->statements.add_raw(n3.asgn);

```

Bison will ultimately pass you the final parsed node, which you then have to put into a `One` edge yourself:

```

auto tree = tree::base::One<program::Program>();
tree.set_raw(n4.prog);
tree->dump();

```

↓

```

Program(
  variables: []
  statements: [
    Assignment(
      lhs: <
        Reference(
          name: x
          target --> !MISSING
        )
      >
      rhs: <
        Literal(
          value: 2
        )
      >
    )
  ]
)

```

Note that when you're using Bison, the enum and its types is largely hidden from you. In the actions the above would look more like this:

```
// Action for a literal:
{ $$ = new value::Literal(std::stoi($1)); }

// Action for a reference:
{ $$ = new value::Reference($1); /* and deallocate $1, assuming it's a char* */ }

// Action for an assignment:
{ $$ = new value::Assignment(); $$->lhs.set_raw($1); $$->rhs.set_raw($3); }
```

and so on.

3.2.1 Syntax error recovery

Some parser generators (like Bison) allow you to specify recovery rules in case of a syntax error, so the parser doesn't just die immediately. This is necessary to emit more than a single error message at a time. To help you store recovery information in the tree for post-mortem analysis, tree-gen has a special `error` directive that you can place inside a node class. For example, if you make a recovery rule that assumes any semicolon encountered after a syntax error produces a broken statement, you can make a node like this:

```
# An erroneous statement.
erroneous_statement {
    error;
}
```

Such nodes behave exactly like any other node, with one exception: they always indicate that they're not well-formed.

```
auto err_stmt = tree::base::make<program::ErroneousStatement>();
ASSERT_RAISES(tree::base::NotWellFormed, err_stmt.check_well_formed());
```

↓

```
tree::base::NotWellFormed exception message: ErroneousStatement error node in tree
```

3.2.2 Line number information (and annotations in general)

When you're doing any kind of serious parsing, you'll want to store as much line number and other contextual information as possible, to make your error messages as clear as possible. You could of course add this information to every node in the tree specification file, perhaps using inheritance to prevent excessive repetition, but this gets annoying fast, especially when it comes to tree-gen's well-formedness rules.

Instead, you should be using annotations for this. Annotations are objects added to nodes without you ever having to declare that you're going to add them. If that sounds like black magic to you in the context of C++, well, that's because it is: internally annotations are stored as a map from `std::type_index` to a C++11 backport of `std::any...` it's complicated. But you don't have to worry about it. What matters is that each and every node generated by tree-gen (or anything else that inherits from `tree::annotatable::Annotatable`) can have zero or one annotation of every C++ type attached to it.

The use of annotations is not limited to metadata (although that is its primary purpose); they also allow bits of code to temporarily attach their own data to tree nodes, without the tree definitions needing to be updated to reflect this. This is especially useful when you're using tree-gen in a library, and you want to let users operate on your trees. After all, they wouldn't be able to modify the tree definition file without forking and recompiling your library.

tree-gen has only one special case for annotations, intended for adding line number information to debug dumps. The `location` directive is used for this:

```
// Source location annotation type. The generated dumper will see if this
// annotation exists for a node, and if so, write it to the debug dump using
// the << stream operator.
location primitives::SourceLocation
```

Without going into too much detail about the Annotatable interface (just look at the API docs for that), here's an example of how it would work.

```
n1.lit->set_annotation(primitives::SourceLocation("test", 1, 5));
n2.ref->set_annotation(primitives::SourceLocation("test", 1, 1));
n3.asgn->set_annotation(primitives::SourceLocation("test", 1, 1));
n4.prog->set_annotation(primitives::SourceLocation("test", 1, 1));
tree->dump();
```

↓

```
Program( # test:1:1
  variables: []
  statements: [
    Assignment( # test:1:1
      lhs: <
        Reference( # test:1:1
          name: x
          target --> !MISSING
        )
      >
      rhs: <
        Literal( # test:1:5
          value: 2
        )
      >
    )
  ]
)
```

Note that we can still use these pointers despite ownership having been passed to the node objects because they refer to the same piece of memory. In practice, though, you would do this during the parser actions, just after constructing them.

When serializing and deserializing, annotations are ignored by default; they can be of any type whatsoever, and C++ can't dynamically introspect which types can be (de)serialized and which can't be, after all. So even though the example `SourceLocation` object extends `tree::annotatable::Serializable`, this doesn't work automatically.

```
auto node = tree::base::make<value::Literal>(2);
node->set_annotation(primitives::SourceLocation("test", 1, 1));
std::string cbor = tree::base::serialize(node);
auto node2 = tree::base::deserialize<value::Literal>(cbor);
ASSERT(!node2->has_annotation<primitives::SourceLocation>());

// However, you *can* register annotations types for serialization and
// deserialization if you want to through the
// `tree::annotatable::serdes_registry` singleton. After that, it will
// work.
tree::annotatable::serdes_registry.add<primitives::SourceLocation>("loc");
cbor = tree::base::serialize(node);
node2 = tree::base::deserialize<value::Literal>(cbor);
ASSERT(node2->has_annotation<primitives::SourceLocation>());
```

Two `add()` methods are provided. The one used here assumes that the type has an appropriate `serialize()` method and an appropriate constructor for deserialization, the other allows you to specify them manually using “`std::function`”s. The name is optional but strongly recommended; if not used, a C++-compiler-specific identifier will be used for the type.

Similar to links, annotations aren’t copied as you might expect by `copy()` and `clone()`. Specifically, annotations are stored as `std::shared_ptr`’s to unknown C++ objects, and therefore copying a node only copies the references to the annotations. To fully clone annotations, you’ll either have to `serialize` and `deserialize` the node they belong to (after registering with `serdes_registry` of course), or `clone` them manually.

3.3 Visitor pattern

In the directory example, we avoided the difficulty of dealing with edges to incomplete types, such as a `One<value::Rvalue>`, aside from a single `as_file()` call halfway through. Spamming such typecasts to, for instance, evaluate an `rvalue` expression during constant propagation or interpreting, is not very scalable. It can also lead to headache down the line if you ever have to add more subclasses, as it’s easy to forget or not bother to check for unknown types when initially writing that code that way. `tree-gen` also generates the necessary classes for the visitor pattern for you to avoid this.

In the visitor pattern, you define a class that implements the appropriate visitor interface, and (in this case) the nodes will call methods on this class depending on their type when you pass the visitor object to them through their `visit()` method. Because the visitor interface classes are also generated by `tree-gen` and have sane default behavior, your code will “by default” be future-proof, in the sense that you’ll get an error if something unexpected happens, rather than it failing silently and maybe crashing down the line.

Two visitor base classes are provided, `Visitor<T=void>` and `RecursiveVisitor`. In either case, you have to override `visit_node()`, which is called when a node is encountered for which no specialization is available; this should just throw an appropriate exception. You can then optionally override any node in the node class hierarchy to provide the actual behavior.

The two classes differ in the default behavior when an unspecialized node is encountered: `Visitor` always defers up the class hierarchy, while `RecursiveVisitor` provides a default implementation for node types with edges, where it simply recursively visits the child nodes depth-first.

The `T` type in `Visitor<T>` refers to the return type for the visit methods. It’s fixed to `void` for `RecursiveVisitor`, because if a node has multiple children, there is nothing sensible to return.

Let’s look at what an expression evaluator looks like for this.

```
class RvalueEvaluator : public value::Visitor<int> {
public:
    int visit_node(value::Node &node) override {
        throw std::runtime_error("unknown node type");
    }

    int visit_literal(value::Literal &node) override {
        return node.value;
    }

    int visit_negate(value::Negate &node) override {
        return -node.oper->visit(*this);
    }

    int visit_add(value::Add &node) override {
        return node.lhs->visit(*this) + node.rhs->visit(*this);
    }
};
```

(continues on next page)

(continued from previous page)

```

}

int visit_sub(value::Sub &node) override {
    return node.lhs->visit(*this) - node.rhs->visit(*this);
}

int visit_mul(value::Mul &node) override {
    return node.lhs->visit(*this) * node.rhs->visit(*this);
}

int visit_div(value::Div &node) override {
    return node.lhs->visit(*this) / node.rhs->visit(*this);
}

int visit_reference(value::Reference &node) override {
    return node.target->value;
}
};

// 2 + (3 * 4) = 14
auto expr = tree::base::make<value::Add>(
    tree::base::make<value::Literal>(2),
    tree::base::make<value::Mul>(
        tree::base::make<value::Literal>(3),
        tree::base::make<value::Literal>(4)
    )
);
RvalueEvaluator eval{};
ASSERT(expr->visit(eval) == 14);

```

If we wouldn't have implemented `visit_node()`, the `RvalueEvaluator eval{};` line would break, as `RvalueEvaluator` would be an abstract class because of it. And while the above may look pretty complete, it is in fact not:

```
ASSERT_RAISES(std::runtime_error, value::ErroneousValue().visit(eval));
```

↓

```
std::runtime_error exception message: unknown node type
```

The visitor pattern is more powerful than recursively calling functions for other reasons as well, because they can be specialized through inheritance, state can be maintained in the class as the tree is traversed (the `dump()` method is actually implemented this way internally), it can be written inline within a function despite the recursive nature, and so on. But it's a lot more verbose than the `as_*()` alternative for simple cases. The choice between visitor pattern and casts is therefore usually a matter of personal preference.

3.4 Complete file listings

3.4.1 value.tree

```

1 // Attach \file docstrings to the generated files for Doxygen.
2 # Implementation for value tree classes.
3 source
4

```

(continues on next page)

(continued from previous page)

```
5 # Header for value tree classes.
6 header
7
8 // Include tree base classes.
9 include "tree-base.hpp"
10 tree_namespace tree::base
11
12 // Include primitive types.
13 include "primitives.hpp"
14 import primitives
15
16 // Initialization function to use to construct default values for the tree base
17 // classes and primitives.
18 initialize_function primitives::initialize
19 serdes_functions primitives::serialize primitives::deserialize
20
21 // Source location annotation type. The generated dumper will see if this
22 // annotation exists for a node, and if so, write it to the debug dump using
23 // the << stream operator.
24 location primitives::SourceLocation
25
26 // Set the namespace for the generated classes and attach a docstring.
27 # Namespace for value tree classes.
28 namespace value
29
30 # Variable instance.
31 variable {
32
33     # The name of the variable.
34     name: primitives::Str;
35
36     # The current value of the variable while interpreting.
37     value: primitives::Int;
38
39 }
40
41 # Toplevel expression node.
42 rvalue {
43
44     # A literal integer.
45     literal {
46
47         # The value of the literal.
48         value: primitives::Int;
49
50     }
51
52     # Unary operators.
53     unop {
54
55         # The operand.
56         oper: One<rvalue>;
57
58         # Negation.
59         negate {}
60
61     }
```

(continues on next page)

```
62
63     # Binary operators.
64     binop {
65
66         # Left-hand side of the expression.
67         lhs: One<rvalue>;
68
69         # Right-hand side of the expression.
70         rhs: One<rvalue>;
71
72         # Addition.
73         add {}
74
75         # Subtraction.
76         sub {}
77
78         # Multiplication.
79         mul {}
80
81         # Division.
82         div {}
83
84     }
85
86     # Toplevel node for assignable expressions.
87     lvalue {
88
89         # A variable reference.
90         reference {
91
92             # The name used to refer to the variable.
93             name: primitives::Str;
94
95             # The variable being referenced.
96             // Note that we use a link here to allow multiple places in the
97             // tree to refer to the same variable instance in our toy
98             // interpreter, so we don't have to do a hashmap lookup each time
99             // and a well-formed tree by definition cannot refer to variables
100             // that don't exist.
101             target: Link<variable>;
102
103         }
104
105         # An erroneous expression.
106         erroneous_value {
107             error;
108         }
109
110     }
111 }
112
```

3.4.2 program.tree

```

1 // Attach \file docstrings to the generated files for Doxygen.
2 # Implementation for program tree classes.
3 source
4
5 # Header for program tree classes.
6 header
7
8 // Include tree base classes.
9 include "tree-base.hpp"
10 tree_namespace tree::base
11
12 // Include primitive types.
13 include "primitives.hpp"
14 import primitives
15
16 // Include the generated value tree header.
17 include "value.hpp"
18 import value
19
20 // Initialization function to use to construct default values for the tree base
21 // classes and primitives.
22 initialize_function primitives::initialize
23 serdes_functions primitives::serialize primitives::deserialize
24
25 // Source location annotation type. The generated dumper will see if this
26 // annotation exists for a node, and if so, write it to the debug dump using
27 // the << stream operator.
28 location primitives::SourceLocation
29
30 // Set the namespace for the generated classes and attach a docstring.
31 # Namespace for program tree classes.
32 namespace program
33
34 # A program is a collection of statements.
35 program {
36
37     # Variable declarations for the program.
38     variables: external Any<value::Variable>;
39
40     # The statements in the program.
41     statements: Any<statement>;
42
43 }
44
45
46 # A statement.
47 statement {
48
49     # Conditional statement.
50     if_else {
51
52         # Branch condition.
53         cond: external One<value::Rvalue>;
54
55         # If block.

```

(continues on next page)

```

56     if_block: Any<statement>;
57
58     # Else block.
59     else_block: Any<statement>;
60
61 }
62
63 # For loop, from start to (non-inclusive) end.
64 for_loop {
65
66     # Variable to assign.
67     var: external One<value::Lvalue>;
68
69     # Start value.
70     start: external One<value::Rvalue>;
71
72     # Stop value.
73     stop: external One<value::Rvalue>;
74
75     # The repeated code.
76     block: Any<statement>;
77
78 }
79
80 # Assignment statement.
81 assignment {
82
83     # Variable to assign.
84     lhs: external One<value::Lvalue>;
85
86     # Expression to assign it to.
87     rhs: external One<value::Rvalue>;
88
89 }
90
91 # Print statement.
92 print {
93
94     # Value to print.
95     expr: external One<value::Rvalue>;
96
97 }
98
99 # An erroneous statement.
100 erroneous_statement {
101     error;
102 }
103
104 }

```

3.4.3 primitives.hpp

```

1  /** \file
2   * Defines primitives used in the generated directory tree structure.
3   */

```

(continues on next page)

(continued from previous page)

```
4
5 #pragma once
6
7 #include <string>
8 #include "tree-base.hpp"
9
10 /**
11  * Namespace with primitives used in the generated directory tree structure.
12  */
13 namespace primitives {
14
15 /**
16  * Integer primitive.
17  */
18 using Int = int;
19
20 /**
21  * Strings, used to represent filenames and file contents.
22  */
23 using Str = std::string;
24
25 /**
26  * Initialization function. This must be specialized for any types used as
27  * primitives in a tree that are actual C primitives (int, char, bool, etc),
28  * as these are not initialized by the T() construct.
29  */
30 template <class T>
31 T initialize() { return T(); };
32
33 /**
34  * Declare the default initializer for integers. It's declared inline to avoid
35  * having to make a cpp file just for this.
36  */
37 template <>
38 inline Int initialize<Int>() {
39     return 0;
40 }
41
42 /**
43  * Serialization function. This must be specialized for any types used as
44  * primitives in a tree. The default implementation doesn't do anything.
45  */
46 template <typename T>
47 void serialize(const T &obj, tree::cbor::MapWriter &map) {
48 }
49
50 /**
51  * Serialization function for Int.
52  */
53 template <>
54 inline void serialize<Int>(const Int &obj, tree::cbor::MapWriter &map) {
55     map.append_int("val", obj);
56 }
57
58 /**
59  * Serialization function for Str.
60  */
```

(continues on next page)

```

61 template <>
62 inline void serialize<Str>(const Str &obj, tree::cbor::MapWriter &map) {
63     map.append_string("val", obj);
64 }
65
66 /**
67  * Deserialization function. This must be specialized for any types used as
68  * primitives in a tree. The default implementation doesn't do anything.
69  */
70 template <typename T>
71 T deserialize(const tree::cbor::MapReader &map) {
72     return initialize<T>();
73 }
74
75 /**
76  * Deserialization function for Int.
77  */
78 template <>
79 inline Int deserialize<Int>(const tree::cbor::MapReader &map) {
80     return map.at("val").as_int();
81 }
82
83 /**
84  * Deserialization function for Str.
85  */
86 template <>
87 inline Str deserialize<Str>(const tree::cbor::MapReader &map) {
88     return map.at("val").as_string();
89 }
90
91 /**
92  * Source location annotation object, containing source file line numbers etc.
93  */
94 class SourceLocation: tree::annotatable::Serializable {
95 public:
96     std::string filename;
97     int line;
98     int column;
99
100     SourceLocation(
101         const std::string &filename,
102         uint32_t line = 0,
103         uint32_t column = 0
104     ) : filename(filename), line(line), column(column) {
105     }
106
107     /**
108      * Serialization logic for SourceLocation.
109      */
110     inline void serialize(tree::cbor::MapWriter &map) const override {
111         map.append_string("filename", filename);
112         map.append_int("line", line);
113         map.append_int("column", column);
114     }
115
116     /**
117      * Deserialization logic for SourceLocation.

```

(continues on next page)

(continued from previous page)

```

118     */
119     inline SourceLocation(const tree::cbor::MapReader &map) {
120         filename = map.at("filename").as_string();
121         line = map.at("line").as_int();
122         column = map.at("column").as_int();
123     }
124
125 };
126
127 /**
128  * Stream << overload for source location objects.
129  */
130 inline std::ostream& operator<<(std::ostream& os, const primitives::SourceLocation&
131 ↪object) {
132     os << object.filename << ":" << object.line << ":" << object.column;
133     return os;
134 }
135 } // namespace primitives

```

3.4.4 CMakeLists.txt

```

1 cmake_minimum_required(VERSION 3.12 FATAL_ERROR)
2
3 project(interpreter-example CXX)
4
5 # Add the tree-gen repository root directory. Normally, your project would be
6 # located outside of tree-gen's root, which means you can omit the second
7 # argument.
8 add_subdirectory(..../.. tree-gen)
9
10 # Generates the files for the value tree.
11 generate_tree_py(
12     "${CMAKE_CURRENT_SOURCE_DIR}/value.tree"
13     "${CMAKE_CURRENT_BINARY_DIR}/value.hpp"
14     "${CMAKE_CURRENT_BINARY_DIR}/value.cpp"
15     "${CMAKE_CURRENT_BINARY_DIR}/value.py"
16 )
17
18 # Generates the files for the program tree.
19 generate_tree_py(
20     "${CMAKE_CURRENT_SOURCE_DIR}/program.tree"
21     "${CMAKE_CURRENT_BINARY_DIR}/program.hpp"
22     "${CMAKE_CURRENT_BINARY_DIR}/program.cpp"
23     "${CMAKE_CURRENT_BINARY_DIR}/program.py"
24 )
25
26 add_executable(
27     interpreter-example
28     "${CMAKE_CURRENT_BINARY_DIR}/value.cpp"
29     "${CMAKE_CURRENT_BINARY_DIR}/program.cpp"
30     "${CMAKE_CURRENT_SOURCE_DIR}/main.cpp"
31 )
32
33 target_include_directories(

```

(continues on next page)

```
34     interpreter-example
35     # This directory for primitives.hpp:
36     PRIVATE "${CMAKE_CURRENT_SOURCE_DIR}"
37     # Binary directory for directory.hpp:
38     PRIVATE "${CMAKE_CURRENT_BINARY_DIR}"
39 )
40
41 target_link_libraries(interpreter-example tree-lib)
42
43 # The following lines only serve to register the example as a test, so you can
44 # run it using `make test` or `ctest` as well. You don't need them as such in
45 # your own project.
46 enable_testing()
47 add_test(interpreter-example interpreter-example)
```

3.4.5 main.cpp

```
1  #include <iostream>
2  #include <cstdio>
3  #include <stdexcept>
4  #include "../utils.hpp"
5
6  // Include the generated files.
7  #include "value.hpp"
8  #include "program.hpp"
9
10 // Note: the // comment contents of main(), together with the MARKER lines and
11 // the output of the program, are used to automatically turn this into a
12 // restructured-text page for ReadTheDocs.
13
14 int main() {
15
16     // *****
17     // Interpreter example
18     // *****
19     //
20     // This example discusses some of the more advanced features of tree-gen,
21     // using an interpreter for a very simple language as an example. This is,
22     // however, not a completely integrated or even functional interpreter; we
23     // just go over some of the concepts and things you may run into when you
24     // would make such an interpreter.
25     //
26     // External nodes & multiple trees
27     // =====
28     //
29     // When trees get large and complicated, you may want to split a tree up
30     // into multiple files. This is supported within tree-gen, as long as there
31     // are no circular references to nodes between the files. In this example,
32     // we've split up the expression-like nodes into a tree separate from the
33     // statement-like nodes. The former, `value.tree`, has nothing special
34     // going on in this regard, while the latter, `program.tree`, effectively
35     // includes the former. The files are listed at the bottom of this page
36     // for your convenience.
37     //
38     // tree-gen itself is unaware of any links between files. For
```

(continues on next page)

(continued from previous page)

```

39 // ``program.tree``, it simply generates code assuming that the types
40 // referred to through the ``external`` edges exist. For this to be true,
41 // you have to tell tree-gen to include the external tree's generated
42 // header file at the top of its generated header file using the ``include``
43 // directive. Note also that, as tree-gen is unaware of the external tree,
44 // you'll have to use the full C++ TitleCase type name, rather than the
45 // snake_case name for normal edges.
46 //
47 // Integration with Flex/Bison (or other parsers)
48 // =====
49 //
50 // When making an interpreter (or parser of any kind) with tree-gen, you'll
51 // probably want to use a tokenizer/parser generator such as Flex/Bison.
52 // With Bison at least, you'll quickly run into a fundamental issue when
53 // trying to store tree-gen nodes in the generated ``yyval`` union:
54 // union in C++ can only consist of types with trivial constructors, and
55 // tree-gen nodes are not among those. The author is unfamiliar with
56 // other frameworks like it, but this will probably apply for any C-style
57 // parser generator.
58 //
59 // The solution is to use raw pointers to the generated node types in this
60 // enumeration. For instance, you may end up with the following (which would
61 // be generated by Bison):
62 union node_type {
63     program::Program *prog;
64     program::Assignment *asgn;
65     value::Literal *lit;
66     value::Reference *ref;
67     // ...
68 };
69 MARKER
70
71 // To create a new node, instead of using ``tree::base::make``,
72 // you use the ``new`` keyword directly.
73 node_type n1, n2;
74 n1.lit = new value::Literal(2);
75 n2.ref = new value::Reference("x");
76 MARKER
77
78 // The generated constructors won't help you much for non-leaf nodes,
79 // because they expect edge wrappers (One, Maybe, etc.) around the
80 // parameters, so you'll typically want to use the constructor without
81 // arguments. Instead, to assign child nodes, you can use
82 // ``Maybe::set_raw()`` , ``One::set_raw()`` , ``Any::add_raw()`` ,
83 // or ``Many::add_raw()`` . These functions take ownership of a new-allocated
84 // raw pointer to a node - exactly what we need here. Note that because they
85 // take ownership, you don't have to (and shouldn't!) delete manually.
86 node_type n3;
87 n3.asgn = new program::Assignment();
88 n3.asgn->rhs.set_raw(n1.lit);
89 n3.asgn->lhs.set_raw(n2.ref);
90
91 node_type n4;
92 n4.prog = new program::Program();
93 n4.prog->statements.add_raw(n3.asgn);
94 MARKER
95

```

(continues on next page)

(continued from previous page)

```

96 // Bison will ultimately pass you the final parsed node, which you then have
97 // to put into a ``One`` edge yourself:
98 auto tree = tree::base::One<program::Program>();
99 tree.set_raw(n4.prog);
100 tree->dump();
101 MARKER
102
103 // Note that when you're using Bison, the enum and its types is largely
104 // hidden from you. In the actions the above would look more like this:
105 //
106 // .. code-block:: none
107 //
108 // // Action for a literal:
109 // { $$ = new value::Literal(std::stoi($1)); }
110 //
111 // // Action for a reference:
112 // { $$ = new value::Reference($1); /* and deallocate $1, assuming it's a char*_
113 ↪ */ }
114 //
115 // // Action for an assignment:
116 // { $$ = new value::Assignment(); $$->lhs.set_raw($1); $$->rhs.set_raw($3); }
117 //
118 // and so on.
119 //
120 // Syntax error recovery
121 // -----
122 //
123 // Some parser generators (like Bison) allow you to specify recovery rules
124 // in case of a syntax error, so the parser doesn't just die immediately.
125 // This is necessary to emit more than a single error message at a time.
126 // To help you store recovery information in the tree for post-mortem
127 // analysis, tree-gen has a special ``error`` directive that you can place
128 // inside a node class. For example, if you make a recovery rule that
129 // assumes any semicolon encountered after a syntax error produces a broken
130 // statement, you can make a node like this:
131 //
132 // .. code-block:: none
133 //
134 // # An erroneous statement.
135 // erroneous_statement {
136 //     error;
137 // }
138 //
139 // Such nodes behave exactly like any other node, with one exception: they
140 // always indicate that they're not well-formed.
141 auto err_stmt = tree::base::make<program::ErroneousStatement>();
142 ASSERT_RAISES(tree::base::NotWellFormed, err_stmt.check_well_formed());
143 MARKER;
144
145 // Line number information (and annotations in general)
146 // -----
147 //
148 // When you're doing any kind of serious parsing, you'll want to store as
149 // much line number and other contextual information as possible, to make
150 // your error messages as clear as possible. You could of course add this
151 // information to every node in the tree specification file, perhaps using
152 // inheritance to prevent excessive repetition, but this gets annoying fast,

```

(continues on next page)

(continued from previous page)

```

152 // especially when it comes to tree-gen's well-formedness rules.
153 //
154 // Instead, you should be using annotations for this. Annotations are
155 // objects added to nodes without you ever having to declare that you're
156 // going to add them. If that sounds like black magic to you in the context
157 // of C++, well, that's because it is: internally annotations are stored
158 // as a map from `std::type_index` to a C++11 backport of `std::any`...
159 // it's complicated. But you don't have to worry about it. What matters is
160 // that each and every node generated by tree-gen (or anything else that
161 // inherits from `tree::annotatable::Annotatable`) can have zero or one
162 // annotation of every C++ type attached to it.
163 //
164 // The use of annotations is not limited to metadata (although that is its
165 // primary purpose); they also allow bits of code to temporarily attach
166 // their own data to tree nodes, without the tree definitions needing to be
167 // updated to reflect this. This is especially useful when you're using
168 // tree-gen in a library, and you want to let users operate on your trees.
169 // After all, they wouldn't be able to modify the tree definition file
170 // without forking and recompiling your library.
171 //
172 // tree-gen has only one special case for annotations, intended for adding
173 // line number information to debug dumps. The `location` directive is
174 // used for this:
175 //
176 // .. code-block:: none
177 //
178 // // Source location annotation type. The generated dumper will see if this
179 // // annotation exists for a node, and if so, write it to the debug dump using
180 // // the << stream operator.
181 // location primitives::SourceLocation
182 //
183 // Without going into too much detail about the Annotatable interface (just
184 // look at the API docs for that), here's an example of how it would work.
185 n1.lit->set_annotation(primitives::SourceLocation("test", 1, 5));
186 n2.ref->set_annotation(primitives::SourceLocation("test", 1, 1));
187 n3.asgn->set_annotation(primitives::SourceLocation("test", 1, 1));
188 n4.prog->set_annotation(primitives::SourceLocation("test", 1, 1));
189 tree->dump();
190 MARKER
191
192 // Note that we can still use these pointers despite ownership having been
193 // passed to the node objects because they refer to the same piece of
194 // memory. In practice, though, you would do this during the parser actions,
195 // just after constructing them.
196 //
197 // When serializing and deserializing, annotations are ignored by default;
198 // they can be of any type whatsoever, and C++ can't dynamically introspect
199 // which types can be (de)serialized and which can't be, after all. So even
200 // though the example SourceLocation object extends
201 // `tree::annotatable::Serializable`, this doesn't work automatically.
202 auto node = tree::base::make<value::Literal>(2);
203 node->set_annotation(primitives::SourceLocation("test", 1, 1));
204 std::string cbor = tree::base::serialize(node);
205 auto node2 = tree::base::deserialize<value::Literal>(cbor);
206 ASSERT(!node2->has_annotation<primitives::SourceLocation>());
207
208 // However, you can register annotations types for serialization and

```

(continues on next page)

```
209 // deserialization if you want to through the
210 // ``tree::annotatable::serdes_registry`` singleton. After that, it will
211 // work.
212 tree::annotatable::serdes_registry.add<primitives::SourceLocation>("loc");
213 cbor = tree::base::serialize(node);
214 node2 = tree::base::deserialize<value::Literal>(cbor);
215 ASSERT(node2->has_annotation<primitives::SourceLocation>());
216 MARKER
217
218 // Two ``add()`` methods are provided. The one used here assumes that the
219 // type has an appropriate ``serialize()`` method and an appropriate
220 // constructor for deserialization, the other allows you to specify them
221 // manually using ``std::function``'s. The name is optional but strongly
222 // recommended; if not used, a C++-compiler-specific identifier will be used
223 // for the type.
224 //
225 // Similar to links, annotations aren't copied as you might expect by
226 // ``copy()`` and ``clone()``. Specifically, annotations are stored as
227 // ``std::shared_ptr``'s to unknown C++ objects, and therefore copying a
228 // node only copies the references to the annotations. To fully clone
229 // annotations, you'll either have to serialize and deserialize the node
230 // they belong to (after registering with ``serdes_registry`` of course), or
231 // clone them manually.
232
233 // Visitor pattern
234 // =====
235 //
236 // In the directory example, we avoided the difficulty of dealing with edges
237 // to incomplete types, such as a ``One<value::Rvalue>``, aside from a
238 // single ``as_file()`` call halfway through. Spamming such typecasts to,
239 // for instance, evaluate an rvalue expression during constant propagation
240 // or interpreting, is not very scalable. It can also lead to headache down
241 // the line if you ever have to add more subclasses, as it's easy to forget
242 // or not bother to check for unknown types when initially writing that code
243 // that way. tree-gen also generates the necessary classes for the visitor
244 // pattern for you to avoid this.
245 //
246 // In the visitor pattern, you define a class that implements the
247 // appropriate visitor interface, and (in this case) the nodes will call
248 // methods on this class depending on their type when you pass the visitor
249 // object to them through their ``visit()`` method. Because the visitor
250 // interface classes are also generated by tree-gen and have sane default
251 // behavior, your code will "by default" be future-proof, in the sense that
252 // you'll get an error if something unexpected happens, rather than it
253 // failing silently and maybe crashing down the line.
254 //
255 // Two visitor base classes are provided, ``Visitor<T=void>`` and
256 // ``RecursiveVisitor``. In either case, you have to override
257 // ``visit_node()`` , which is called when a node is encountered for which
258 // no specialization is available; this should just throw an appropriate
259 // exception. You can then optionally override any node in the node class
260 // hierarchy to provide the actual behavior.
261 //
262 // The two classes differ in the default behavior when an unspecialized
263 // node is encountered: ``Visitor`` always defers up the class hierarchy,
264 // while ``RecursiveVisitor`` provides a default implementation for node
265 // types with edges, where it simply recursively visits the child nodes
```

(continues on next page)

(continued from previous page)

```

266 // depth-first.
267 //
268 // The T type in ``Visitor<T>`` refers to the return type for the visit
269 // methods. It's fixed to ``void`` for ``RecursiveVisitor``, because if a
270 // node has multiple children, there is nothing sensible to return.
271 //
272 // Let's look at what an expression evaluator looks like for this.
273 class RvalueEvaluator : public value::Visitor<int> {
274 public:
275     int visit_node(value::Node &node) override {
276         throw std::runtime_error("unknown node type");
277     }
278
279     int visit_literal(value::Literal &node) override {
280         return node.value;
281     }
282
283     int visit_negate(value::Negate &node) override {
284         return -node.oper->visit(*this);
285     }
286
287     int visit_add(value::Add &node) override {
288         return node.lhs->visit(*this) + node.rhs->visit(*this);
289     }
290
291     int visit_sub(value::Sub &node) override {
292         return node.lhs->visit(*this) - node.rhs->visit(*this);
293     }
294
295     int visit_mul(value::Mul &node) override {
296         return node.lhs->visit(*this) * node.rhs->visit(*this);
297     }
298
299     int visit_div(value::Div &node) override {
300         return node.lhs->visit(*this) / node.rhs->visit(*this);
301     }
302
303     int visit_reference(value::Reference &node) override {
304         return node.target->value;
305     }
306 };
307
308 // 2 + (3 * 4) = 14
309 auto expr = tree::base::make<value::Add>(
310     tree::base::make<value::Literal>(2),
311     tree::base::make<value::Mul>(
312         tree::base::make<value::Literal>(3),
313         tree::base::make<value::Literal>(4)
314     )
315 );
316 RvalueEvaluator eval{};
317 ASSERT(expr->visit(eval) == 14);
318 MARKER
319
320 // If we wouldn't have implemented ``visit_node()`` , the
321 // ``RvalueEvaluator eval{};`` line would break, as RvalueEvaluator would
322 // be an abstract class because of it. And while the above may look pretty

```

(continues on next page)

(continued from previous page)

```
323 // complete, it is in fact not:
324 ASSERT_RAISES(std::runtime_error, value::ErroneousValue().visit(eval));
325 MARKER
326
327 // The visitor pattern is more powerful than recursively calling functions
328 // for other reasons as well, because they can be specialized through
329 // inheritance, state can be maintained in the class as the tree is
330 // traversed (the ``dump()`` method is actually implemented this way
331 // internally), it can be written inline within a function despite the
332 // recursive nature, and so on. But it's a lot more verbose than the
333 // ``as_*()`` alternative for simple cases. The choice between visitor
334 // pattern and casts is therefore usually a matter of personal preference.
335 MARKER
336
337 return 0;
338 }
```